arXiv:1410.0719v2 [cs.NA] 9 Oct 2014

Proceedings of the second "international Traveling Workshop on Interactions between Sparse models and Technology" iTWIST'14

L'Arsenal, Namur, Belgium

August 27-29, 2014.

iTWIST'14 Group Photo

iTWIST'14 Presentation

The second edition of the "*international - Traveling Workshop on Interactions between Sparse models and Technology*" (iTWIST) took place in the medieval and picturesque town of Namur in Belgium. The workshop was conveniently located in "The Arsenal" building within walking distance of both hotels and town center. One implicit objective of this biennial workshop is to foster collaboration between international scientific teams by disseminating ideas through both specific oral/poster presentations and free discussions.
For this second edition, iTWIST'14 has gathered about 70 international participants and has featured 9 invited talks, 10 oral presentations, and 14 posters on the following themes, all related to the theory, application and generalization of the "sparsity paradigm":

- Sparsity-driven data sensing and processing
- Beyond linear and convex inverse problem
- Blind inverse problems and dictionary learning
- Information theory, geometry and randomness
- Sparsity? What's next?
- Union of low dimensional subspaces
- Matrix/manifold/graph sensing/processing
- Sparsity and computational neuroscience
- Complexity/accuracy tradeoffs in numerical methods
- Sparse machine learning and inference

Pictures of "L'Arsenal" Rue de Bruno, 11 5000 Namur - Belgique

Scientific Organizing Committee:

- Laurent Jacques (UCL, Belgium)
- Yannick Boursier (CPPM, Aix-Marseille U., France)
- Christine De Mol (ULB, Belgium)
- Sandrine Anthoine (I2M, Aix-Marseille U., France)
- Pascal Frossard (EPFL, Switzerland).
- Christophe De Vleeschouwer (UCL, Belgium)
- Prasad Sudhakar (UCL, Belgium)
- Aleksandra Pizurica (Ghent U., Belgium)
- Pierre Vandergheynst (EPFL, Switzerland)

Local Organizing Committee:

- Laurent Jacques (UCL, Belgium)
- Jean Deschuyter (UCL, Belgium)
- Kévin Degraux (UCL, Belgium)
- Pascaline Parisot (UCL, Belgium)
- Muhammad Arsalan (UCL, Belgium)
- Amit Kumar KC (UCL, Belgium)
- Christophe De Vleeschouwer (UCL, Belgium)
- Adriana Gonzalez (UCL, Belgium)
- Stéphanie Guérit (UCL, Belgium)
- Nicolas Matz (INSA, Toulouse, France)
- Cédric Verleysen (UCL, Belgium)

Sponsors: The iTWIST'14 organizing committee thanks the following sponsors for their help and fundings.

Table of contents

(p. 1) C. Bilen (INRIA Rennes, France), Srdan Kitic (INRIA Rennes, France), N. Bertin (IRISA CNRS, France), R. Gribonval (INRIA Rennes, France),
"Sparse Acoustic Source Localization with Blind Calibration for Unknown Medium Characteristics".

(p. 4) N. Boumal (UCL, Belgium), B. Mishra (ULg, Belgium), P.-A. Absil (UCL, Belgium), R. Sepulchre (Cambridge U., UK),
"Manopt: a Matlab toolbox for optimization on manifolds".

(p. 7) S. Bundervoet (VUB-ETRO,Belgium), C. Schretter (VUB-ETRO,Belgium), A.Dooms (VUB-ETRO,Belgium), P. Schelkens (VUB-ETRO, Belgium),
"Bayesian Estimation of Sparse Smooth Speckle Shape Models for Motion Tracking in Medical Ultrasound".

(p. 10) O. Chabiron (U.Toulouse,France), F.Malgouyres (U.Toulouse,France), J.-Y.,Tourneret (U.Toulouse,France), N. Dobigeon (U. Toulouse, France),
"Learning a fast transform with a dictionary".

(p. 13) P. Chainais (LAGIS / INRIA Lille, France), C. Richard (Lab. Lagrange, Nice, France),
"A diffusion strategy for distributed dictionary learning".

(p. 16) B. Cornelis (VUB-ETRO, Belgium), A. Dooms (VUB-ETRO, Belgium), I. Daubechies (Duke U., USA), D. Dunson (Duke U., USA),
"Bayesian crack detection in high resolution data".

(p. 19) M. Dankova (Brno University of Technology) and P. Rajmic (Brno University of Technology, Czech Republic)
"Compressed sensing of perfusion MRI".

(p. 21) K. Degraux (UCL, Belgium), V. Cambareri (U.Bologna, Italy), B.Geelen (IMEC, Belgium), L. Jacques (UCL, Belgium), G. Lafruit (IMEC, Belgium), G. Setti (U. Bologna, Italy),
"Compressive Hyperspectral Imaging by Out-of-Focus Modulations and Fabry-Pérot Spectral Filters".

(p. 24) J.-F. Determe (ULB, Belgium), J. Louveaux (UCL, Belgium), F. Horlin (ULB, Belgium),
"Filtered Orthogonal Matching Pursuit: Applications".

(p. 26) A. Drémeau (ESPCI ParisTech, France), P. Héas (INRIA, Rennes, France), C. Herzet (INRIA, Rennes, France),
"Combining sparsity and dynamics: an efficient way".

(p. 29) V. Duval (U. Paris-Dauphine), G. Peyré (U. Paris-Dauphine),
"Discrete vs. Continuous Sparse Regularization".

(p. 31) A. Fawzi (LTS4, EPFL, Switzerland), M. Davies (U. Edinburgh, UK), P. Frossard (LTS4, EPFL, Switzerland),
"Dictionary learning for efficient classification based on soft-thresholding".

(p. 34) N. Gillis (U. Mons, Belgium), S. A. Vavasis (U. Waterloo, UK),
"Semidefinite Programming Based Preconditioning for More Robust Near-Separable Nonnegative Matrix Factorization".

(p. 36) C. Herzet (INRIA Rennes, France), C. Soussen (U. de Lorraine, Nancy, France),
"Enhanced Recovery Conditions for OMP/OLS by Exploiting both Coherence and Decay".

(p. 38) S. Kitic, (Inria, France), N. Bertin, (IRISA, France), R. Gribonval (Inria, France),
"Wideband Audio Declipping by Cosparse Hard Thresholding".

(p. 41) L. Le Magoarou (INRIA, Rennes, France), R. Gribonval (INRIA, Rennes, France),
"Strategies to learn computationally efficient and compact dictionaries".

(p. 44) J. Liang (GREYC, ENSICAEN, France), J. Fadili (GREYC, ENSICAEN, France), G. Peyré (U. Paris-Dauphine, France),
"Iteration-Complexity for Inexact Proximal Splitting Algorithms".

(p. 46) A. Liutkus (Institut Langevin, France), D. Martina (Institut Langevin, France), S. Gigan (Institut Langevin, France), L. Daudet (Institut Langevin, France),
"Calibration and imaging through scattering media".

(p. 48) M. Maggioni (Duke U.), S. Minsker (Duke U.), N. Strawn (Duke U.),
"Multiscale Dictionary and Manifold Learning: Non-Asymptotic Bounds for the Geometric Multi-Resolution Analysis".

(p. 51) C. Mory (iMagX Belgium and CREATIS U. Lyon France) and L. Jacques (UCLouvain, Belgium)
"An application of the Chambolle-Pock algorithm to 3D + time tomography".

(p. 54) F. Ngole (CEA-CNRS-Paris 7, France), Starck, Jean-Luc (CEA-CNRS-Paris 7, France),
"Super-resolution method using sparse regularization for point spread function recovery".

(p. 57) C. Schretter (VUB-ETRO, Belgium), I. Loris (ULB, Belgium), A. Doom (VUB-ETRO, Belgium), P. Schelkens (VUB-ETRO, Belgium),
"Total Variation Reconstruction From Quasi-Random Samples".

(p. 59) S. Vaiter (U. Paris-Dauphine, France), M. Golbabaee (U. Paris-Dauphine, France), J. Fadili (GREYC, ENSICAEN, France), G. Peyré (U. Paris-Dauphine, France),
"Model Selection with Piecewise Regular Gauges".

(p. 61) D. Vukobratovic (U. of Novi Sad, Serbia), A. Pizurica (Ghent University, Belgium),
"Adaptive Compressed Sensing Using Sparse Measurement Matrices".

Sparse Acoustic Source Localization with Blind Calibration for Unknown Medium Characteristics

Çağdaş Bilen[1], Srđan Kitić[1], Nancy Bertin[2] and Rémi Gribonval[1]

[1]INRIA, Centre Inria Rennes - Bretagne Atlantique, 35042 Rennes Cedex, France.*

[2]IRISA - CNRS UMR 6074, Campus de Beaulieu, F-35042 Rennes Cedex, Rennes, France.

***Abstract*— We consider the problem of audio source localization in a medium with unknown characteristics, particularly the speed of sound within the medium. We propose an algorithm to retrieve both the speed of sound and the localized signals under reasonable assumptions such as smoothness of medium structure. We present initial simulation results for stationary sinusoidal sources both for the case of uniform and non-uniform speed of sound through the medium to demonstrate the performance of the proposed algorithm.**

1 Introduction

Sparse recovery has become a popular research field especially after the emergence of compressed sensing theory [1]. Consequently many natural signals have also been shown to be compressible or essentially low dimensional which enabled sparsity inducing methods to be used for the estimation of these signals in linear inverse problems.

Acoustic source localization can also be formulated as a sparse recovery problem [2, 3]. In the continuous space and time, the acoustic pressure field in a 2-dimensional space, $p(\overrightarrow{s},t)$, induced by the pressure emitted by the sources, $f(\overrightarrow{s},t)$, is known to obey the differential wave equation

$$\Delta p(\overrightarrow{s},t) - \frac{1}{c^2}\frac{\partial^2 p(\overrightarrow{s},t)}{\partial t^2} = \begin{cases} 0, & \text{if no source at } \overrightarrow{s} \\ f(\overrightarrow{s},t), & \text{if source at } \overrightarrow{s} \end{cases} \quad (1)$$

at all non-boundary locations. The parameter c in (1) represents the speed of sound in the medium. For inhomogeneous mediums, this constant can be represented more generally as $c(\overrightarrow{s},t)$. In addition to the wave equation, the acoustic pressure field is also restricted by the boundary conditions. Both the wave equation and the boundary equations can be discretized straightforwardly into a set of linear differential equations represented as

$$(\mathbf{\Omega}_1 + \alpha\mathbf{\Omega}_2)\mathbf{x} = \mathbf{z} \quad (2)$$

in which $\mathbf{\Omega}_1$ and $\mathbf{\Omega}_2$ represent the discretized differential operators in space and time respectively, $\alpha \triangleq 1/c^2$ is a function of speed of sound, $\mathbf{x}$ and $\mathbf{z}$ represent the discretized pressure field and the source signals in vector form respectively.

The source localization problem can be described as the estimation of $\mathbf{z}$ (or equivalently $\mathbf{x}$) from limited number of measurements recorded by a number of microphones, i.e. from

$$\mathbf{y} = \mathbf{I}_s\mathbf{x} \quad (3)$$

where $\mathbf{I}_s$ is a matrix formed by keeping some of the rows of the identity matrix corresponding to microphone locations. Considering the sparse nature of the sources, source localization can be posed as a sparse recovery problem which can be solved with convex optimization such as in

$$\hat{\mathbf{x}} = \arg\min_{\mathbf{x}} \|(\mathbf{\Omega}_1 + \alpha\mathbf{\Omega}_2)\mathbf{x}\|_* \text{ s.t. } \mathbf{I}_s\mathbf{x} = \mathbf{y} \quad (4)$$

where the $*$ norm is often chosen as the ℓ_1 norm or the $\ell_{1,2}$ norm [3]. In many such approaches, the medium is assumed to be homogeneous and the medium properties (c) are assumed to be known. However, in practice this assumption may not be always valid since the behavior of the medium may vary depending on the physical conditions such as temperature, ambient pressure, or unknown inhomogeneities in the medium. These incorrect assumptions on α (and hence the sparsifying operator $\mathbf{\Omega}_1 + \alpha\mathbf{\Omega}_2$) may lead to significantly decreased sparsity or multiplicative perturbations that would highly reduce the recovery performance [4, 5].

In this abstract, we investigate the source localization problem formulated in (4) but with an unknown speed of sound c (or equivalently α). In Section 2, we formulate the source localization problem for a spatially varying c, i.e. a time invariant but inhomogeneous medium and present an algorithm to estimate the medium properties along with the sources under spatial smoothness assumptions. Preliminary experimental results with the proposed algorithm are then presented in Section 3, which is followed by the concluding remarks in Section 4.

2 Source Localization with Unknown Medium Parameters

Let us first assume that the medium is inhomogeneous and unknown, i.e. $\alpha = \alpha(\overrightarrow{s},t)$ for the continuous domain. Consequently the discretized equation in (2) can be modified as

$$(\mathbf{\Omega}_1 + \mathrm{Diag}(\boldsymbol{\alpha})\mathbf{\Omega}_2)\mathbf{x} = \mathbf{z} \quad (5)$$

where $\mathrm{Diag}(\boldsymbol{\alpha})$ indicates a diagonal matrix with the entries of the vector $\boldsymbol{\alpha}$ along the diagonal. In this setting one can attempt to recover $\boldsymbol{\alpha}$ and $\mathbf{x}$ by minimizing

$$\hat{\mathbf{x}}, \hat{\boldsymbol{\alpha}} = \arg\min_{\mathbf{x},\boldsymbol{\alpha}} \|(\mathbf{\Omega_1} + \mathrm{Diag}(\boldsymbol{\alpha})\mathbf{\Omega_2})\mathbf{x}\|_{1,2} \text{ s.t. } \mathbf{I}_s\mathbf{x} = \mathbf{y} \quad (6)$$

where the $\ell_{1,2}$ norm is defined as ℓ_1 norm along the spatial and ℓ_2 norm along the temporal dimension, assuming all the source locations are fixed along time. However the global minimum of the objective function in (5) often does not correspond to the actual $\mathbf{x}$ and $\boldsymbol{\alpha}$ due to the high number of degrees of freedom introduced by the unknown $\boldsymbol{\alpha}$.

*This work was partly funded by the European Research Council, PLEASE project (ERC-StG-2011-277906).

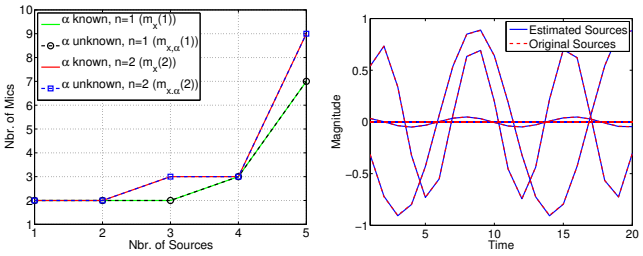

(a) (b)

Figure 1: (a) The minimum sufficient number of the microphones, $m_{\mathbf{x},\boldsymbol{\alpha}}(n)$, with $n = 1, 2$ for successfully localizing different number of sources. The required number of microphones for the ideal case when the field $\boldsymbol{\alpha}$ is known, $m_{\mathbf{x}}(n)$, is also plotted for comparison. (b) The original and estimated signals emitted from the sources for a sample recovery.

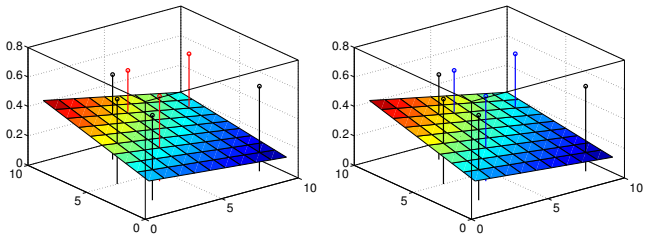

(a) The original field $\boldsymbol{\alpha}$ (b) The estimated field $\hat{\boldsymbol{\alpha}}$

Figure 2: The original field $\boldsymbol{\alpha}$ and the estimated field $\hat{\boldsymbol{\alpha}}$ on a 10 by 10 2-dimensional grid for a sample experiment. The red, blue and black markers indicate the spatial positions of the original sources, estimated sources and the microphones respectively.

The intrinsic degrees of freedom in the structure of a real world medium is often much smaller. In this work we will assume the medium does not change in time and exploit the fact that the spatial variation in a natural medium is often limited and assume nth order spatial difference of the field $\boldsymbol{\alpha}$ along each spatial dimension is zero, i.e.

$$\mathbf{D}_n \boldsymbol{\alpha} = 0 \tag{7}$$

$$\Rightarrow \boldsymbol{\alpha} = \mathbf{A}_n \mathbf{p} \tag{8}$$

where $\mathbf{A}_n$ is an orthogonal basis for the null space of $\mathbf{D}_n$. It can be observed that the degrees of freedom for $\boldsymbol{\alpha}$ are reduced to n^2 (for 2-dimensional space) from the total size of the discretized points in space, which can be a significant reduction in practice. With the assumption of (8), the optimization in (6) can be modified as

$$\hat{\mathbf{x}}, \hat{\mathbf{p}} = \arg\min_{\mathbf{x},\mathbf{p}} \|(\boldsymbol{\Omega}_1 + \mathrm{Diag}(\mathbf{A}_n \mathbf{p})\boldsymbol{\Omega}_2)\mathbf{x}\|_{1,2} \text{ s.t. } \mathbf{I}_s \mathbf{x} = \mathbf{y} \tag{9}$$

Unlike the case with known α, (9) is not a linear optimization when jointly considering the variables $\mathbf{x}$ and α, hence the well developed convex optimization methods are not applicable. It is, however, linear in terms of each of the unknowns when the other is considered a constant. Therefore we propose a form of the alternating minimization approach which is often employed in such scenarios.

3 Experimental Results

In order to demonstrate the joint estimation performance of the proposed approach, a rectangular grid of 10 by 10 spatial dimensions is simulated with k randomly located stationary sources, k varying from 1 to 5. The sources emit random sinusoidals for a duration of 20 time instants. The pressure field is measured by m randomly located stationary microphones for 100 time instants, m varying from 2 to 10. The boundary conditions are modeled to be Neumann boundaries as in [2, 3]. The field $\boldsymbol{\alpha}$ is generated randomly to obey (8) with parameter n set as 1 and 2 and the field values to be between 0.1 and 0.6.

The field parameters $\mathbf{p}$ and the acoustic pressure $\mathbf{x}$ are estimated from the microphone measurements $\mathbf{y} = \mathbf{I}_s \mathbf{x}$ with an algorithm that minimizes (9) by alternating between unknowns applying the method of Alternating Direction Method of Multipliers (ADMM [6]). The details of the algorithm are not provided here due to space constraints.

Each of the randomly generated experiments is repeated 20 times with the same parameters to empirically estimate the minimum sufficient number of microphones for successful source localization for each set of n, m, k. After each simulation, k points with highest energy are chosen as estimates of the source locations. The minimum number of microphones that succeeds in localization of the sources in 80% of the 20 experiments is set as the minimum sufficient number of microphones, $m_{\mathbf{x},\boldsymbol{\alpha}}(n)$. The minimum sufficient number of the microphones for the recovery of the acoustic field when the medium parameter, $\boldsymbol{\alpha}$, is known, $m_{\mathbf{x}}$, is also computed the same way for comparison. Figure 1a shows the change of $m_{\mathbf{x},\boldsymbol{\alpha}}(n)$ and $m_{\mathbf{x}}(n)$ with respect to the number of sources. A sample reconstruction of source signals is shown in Figure 1b. The original and the estimated field $\boldsymbol{\alpha}$ as well as the source and the microphone locations for the same sample reconstruction can also be seen in Figure 2.

The first thing to notice in the presented results in Figure 1a is that estimating the unknown α introduces almost no performance degradation with respect to perfectly knowing the medium parameters. This can be attributed to the very small degrees of freedom in $\boldsymbol{\alpha}$ thanks to the smoothness constraints. Hence estimating $\boldsymbol{\alpha}$ while localizing sources is possible whenever the sources can be localized with known $\boldsymbol{\alpha}$. In fact it is observed in the experiments that the estimation of $\boldsymbol{\alpha}$ is still accurate for many cases when source localization is not successful. A second observation is that the sufficient number of microphones is not significantly affected when $\boldsymbol{\alpha}$ is spatially varying ($n = 2$) compared to when it is spatially constant ($n = 1$). However the degrees of freedom within $\boldsymbol{\alpha}$ increases with n^2 and therefore the sufficient number of microphones is expected to increase more rapidly with increasing n.

4 Conclusion

In this work, the acoustic source localization problem with additional variables due to unknown medium properties is described and a sparse recovery approach is presented for estimating both the sources and the medium structure. The presented approach utilizes the assumption that the spatial variation in the medium is limited and the acoustic sources are sparse. Preliminary results that demonstrate the performance of the proposed approach are presented. The talk will include further results, the details on the recovery algorithm and discussions on the extensions of the algorithm for more realistic scenarios.

References

[1] D. Donoho, "Compressed sensing", Information Theory, IEEE Transactions on, vol. 52, no. 4, 2006.

[2] S. Nam and R. Gribonval, "Physics-driven structured cosparse modeling for source localization", in Acoustics, Speech and Signal Processing (ICASSP), 2012 IEEE International Conference on.

[3] S. Kitić, N. Bertin and R. Gribonval, "A review of cosparse signal recovery methods applied to sound source localization", Le XXIVe colloque Gretsi, Brest, France, Sep 2013.

[4] E. J. Candès, J. K. Romberg, and T. Tao, "Stable signal recovery from incomplete and inaccurate measurements", Communications on Pure and Applied Mathematics, vol. 59, no. 8, pp. 1207-1223, 2006.

[5] M. Herman and T. Strohmer, "General deviants: An analysis of perturbations in compressed sensing", Selected Topics in Signal Processing, IEEE Journal of, vol. 4, pp. 342-349, April 2010.

[6] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein, "Distributed Optimization and Statistical Learning via the Alternating Direction Method of Multipliers", Foundations and Trends in Machine Learning, vol. 3, no. 1, pp. 1-122, 2011.

Manopt: a Matlab toolbox for optimization on manifolds

N. Boumal[1], B. Mishra[2], P.-A. Absil[1] and R. Sepulchre[3].

[1]Université catholique de Louvain, Belgium. [2]Université de Liège, Belgium. [3]University of Cambridge, UK.

***Abstract*— Optimization on manifolds is a rapidly developing branch of nonlinear optimization. Its focus is on problems where the smooth geometry of the search space can be leveraged to design efficient numerical algorithms. In particular, optimization on manifolds is well-suited to deal with rank and orthogonality constraints. Such structured constraints appear pervasively in machine learning applications, including low-rank matrix completion, sensor network localization, camera network registration, independent component analysis, metric learning, dimensionality reduction and so on.**

The Manopt toolbox, available at `www.manopt.org`, is a user-friendly, documented piece of software dedicated to simplify experimenting with state of the art Riemannian optimization algorithms.

This abstract is lightly adapted from (Boumal et al., 2014).

1 Introduction

Optimization on manifolds, or Riemannian optimization, is a fast growing research topic in the field of nonlinear optimization. Its purpose is to provide efficient numerical algorithms to solve optimization problems of the form

$$\min_{x \in \mathcal{M}} f(x), \tag{1}$$

where the search space $\mathcal{M}$ is a smooth space: a differentiable manifold which can be endowed with a Riemannian structure. In a nutshell, this means $\mathcal{M}$ can be linearized locally at each point x as a tangent space $\mathrm{T}_x\mathcal{M}$ and an inner product $\langle \cdot, \cdot \rangle_x$ which smoothly depends on x is available on $\mathrm{T}_x\mathcal{M}$. A number of smooth search spaces arise often in applications:

- The **oblique manifold** is a product of spheres, $\mathcal{M} = \{X \in \mathbb{R}^{n \times m} \colon \mathrm{diag}(X^\top X) = \mathbb{1}_m\}$. That is, $X \in \mathcal{M}$ if each column of X has unit 2-norm in $\mathbb{R}^n$. Absil and Gallivan (2006) show how independent component analysis can be cast on this manifold as non-orthogonal joint diagonalization. When furthermore it is only the product $Y = X^\top X$ which matters, matrices of the form QX are equivalent for all orthogonal Q. Quotienting out this equivalence relation yields the **fixed-rank elliptope** $\mathcal{M} = \{Y \in \mathbb{R}^{m \times m} : Y = Y^\top \succeq 0, \mathrm{rank}(Y) = n, \mathrm{diag}(Y) = \mathbb{1}_m\}$. See the example below for application to the max-cut problem. The packing problem on the sphere, where one wishes to place m points on the unit sphere in $\mathbb{R}^n$ such that the two closest points are as far apart as possible (Dirr et al., 2007), is another example of an optimization problem on the fixed-rank elliptope. Grubišić and Pietersz (2007) optimize over this set to produce low-rank approximations of covariance matrices.

- The (compact) **Stiefel manifold** is the set of orthonormal matrices, $\mathcal{M} = \{X \in \mathbb{R}^{n \times m} \colon X^\top X = I_m\}$. Amari (1999) and Theis et al. (2009) formulate versions of independent component analysis with dimensionality reduction as optimization over the Stiefel manifold.

- The **Grassmann manifold** $\mathcal{M} = \{\mathrm{col}(X) \colon X \in \mathbb{R}_*^{n \times m}\}$ is the manifold, where $\mathbb{R}_*^{n \times m}$ is the set of full-rank matrices in $\mathbb{R}^{n \times m}$ and $\mathrm{col}(X)$ denotes the subspace spanned by the columns of X. That is, $\mathrm{col}(X) \in \mathcal{M}$ is a subspace of $\mathbb{R}^n$ of dimension m. Among other things, optimization over the Grassmann manifold proves useful in low-rank matrix completion, where it is observed that if one knows the column space spanned by the sought matrix, then completing the matrix according to a least squares criterion is easy (Keshavan et al., 2010; Boumal and Absil, 2011; Balzano et al., 2010).

- The **special orthogonal group** is the group of rotations $\mathcal{M} = \{X \in \mathbb{R}^{n \times n} \colon X^\top X = I_n \text{ and } \det(X) = 1\}$, typically considered as a Riemannian submanifold of $\mathbb{R}^{n \times n}$. Optimization problems involving rotation matrices notably occur in robotics and computer vision, when estimating the attitude of vehicles or the pose of cameras (Tron and Vidal, 2009; Boumal et al., 2013).

- **Fixed-rank matrices** $\mathcal{M} = \{X \in \mathbb{R}^{n \times m} \colon \mathrm{rank}(X) = k\}$ admit a number of different Riemannian structures. Vandereycken (2013) proposes an embedded geometry for $\mathcal{M}$ and exploits Riemannian optimization on that manifold to address the low-rank matrix completion problem. Shalit et al. (2012) use the same geometry to address similarity learning. Mishra et al. (2012) cover a number of quotient geometries for $\mathcal{M}$ and similarly address low-rank matrix completion problems.

- The set of **symmetric, positive semidefinite, fixed-rank matrices** is also a manifold, $\mathcal{M} = \{X \in \mathbb{R}^{n \times n} \colon X = X^\top \succeq 0, \mathrm{rank}(X) = k\}$. Meyer et al. (2011) exploit this to propose low-rank algorithms for metric learning. This space is tightly related to the space of **Euclidean distance matrices** X such that X_{ij} is the squared distance between two fixed points $x_i, x_j \in \mathbb{R}^k$. Mishra et al. (2011) leverage this geometry to formulate efficient low-rank algorithms for Euclidean distance matrix completion.

- The **fixed-rank spectrahedron** $\mathcal{M} = \{X \in \mathbb{R}^{n \times n} \colon X = X^\top \succeq 0, \mathrm{trace}(X) = 1 \text{ and } \mathrm{rank}(X) = k\}$, without the rank constraint, is a convex set which can be used to solve relaxed (lifted) formulations of the sparse PCA problem. Journée et al. (2010) show how optimizing over the fixed-rank spectrahedron can lead to efficient algorithms for sparse PCA.

The rich geometry of Riemannian manifolds $\mathcal{M}$ makes it possible to define gradients and Hessians of cost functions f, as well as systematic procedures (called *retractions*) to move on the manifold starting at a point x, along a specified tangent direction at x. Those are sufficient ingredients to generalize standard nonlinear optimization methods such as gradient descent, conjugate-gradients, quasi-Newton, trust-regions, etc.

In a recent monograph, Absil et al. (2008) lay down a mature framework to analyze problems of the form (1) when f is a smooth function, with a strong emphasis on building a theory that leads to efficient numerical algorithms. In particular, they describe the necessary ingredients to design first- and second-order algorithms on Riemannian manifolds in general. These algorithms come with convergence guarantees essentially matching those of the Euclidean counterparts they generalize. For example, the Riemannian trust-region method is known to converge globally toward critical points and to converge locally quadratically when the Hessian of f is available. In many respects, this theory subsumes well-known results from an earlier paper by Edelman et al. (1998), which focused on problems of the form (1) with $\mathcal{M}$ either the set of orthonormal matrices (the Stiefel manifold) or the set of linear subspaces (the Grassmann manifold).

The maturity of the theory of smooth Riemannian optimization, its widespread applicability and its excellent track record performance-wise prompted us to build the Manopt toolbox: a user-friendly piece of software to help researchers and practitioners experiment with these tools. Code and documentation are available at `www.manopt.org`.

2 Architecture and features of Manopt

The toolbox architecture is based on a separation of the manifolds, the solvers and the problem descriptions. For basic use, one only needs to pick a manifold from the library, describe the cost function (and possible derivatives) on this manifold and pass it on to a solver. Accompanying tools help the user in common tasks such as numerically checking whether the cost function agrees with its derivatives up to the appropriate order, approximating the Hessian based on the gradient of the cost, etc.

Manifolds in Manopt are represented as structures and are obtained by calling a factory. The manifold descriptions include projections on tangent spaces, retractions, helpers to convert Euclidean derivatives (gradient and Hessian) to Riemannian derivatives, etc. All the manifolds mentioned in the introduction work out of the box, and more can be added (shape space (Ring and Wirth, 2012), low-rank tensors (Kressner et al., 2013), etc.). Cartesian products of known manifolds are supported too.

Solvers are functions in Manopt that implement generic Riemannian minimization algorithms. Solvers log standard information at each iteration and comply with standard stopping critera. Extra information can be logged via callbacks and, similarly, user-defined stopping criteria are allowed. Currently, Riemannian trust-regions (based on (Absil et al., 2007)) and conjugate-gradients are implemented (with preconditioning), as well as steepest-descent and a couple derivative free schemes. More solvers can be added, with an outlook toward Riemannian BFGS (Ring and Wirth, 2012), stochastic gradients (Bonnabel, 2013), nonsmooth subgradients schemes (Dirr et al., 2007), etc.

An optimization problem in Manopt is represented as a problem structure. The latter includes a field which contains a structure describing a manifold, as obtained from a factory. Additionally, the problem structure hosts function handles for the cost function f and (possibly) its derivatives. An abstraction layer at the interface between the solvers and the problem description offers great flexibility in the cost function description. As the needs grow during the life-cycle of the toolbox and new ways of describing f become necessary (subdifferentials, partial gradients, etc.), it will be sufficient to update this interface.

Computing $f(x)$ typically produces intermediate results which can be reused in order to compute the derivatives of f at x. To prevent redundant computations, Manopt incorporates an (optional) caching system, which becomes useful when transiting from a proof-of-concept draft of the algorithm to a convincing implementation.

3 Example: the maximum cut problem

Given an undirected graph with n nodes and weights $w_{ij} \geq 0$ on the edges such that $W \in \mathbb{R}^{n\times n}$ is the weighted adjacency matrix and $D \in \mathbb{R}^{n\times n}$ is the diagonal degree matrix with $D_{ii} = \sum_j w_{ij}$, the graph Laplacian is the positive semidefinite matrix $L = D - W$. The max-cut problem consists in building a partition $s \in \{+1, -1\}^n$ of the nodes in two classes such that $\frac{1}{4}s^\top Ls = \sum_{i<j} w_{ij} \frac{(s_i - s_j)^2}{4}$, i.e., the sum of the weights of the edges connecting the two classes, is maximum. Let $X = ss^\top$. Then, max-cut is equivalent to:

$$\max_{X\in\mathbb{R}^{n\times n}} \operatorname{trace}(LX)/4$$
$$\text{s.t. } X = X^\top \succeq 0, \operatorname{diag}(X) = \mathbb{1}_n \text{ and } \operatorname{rank}(X) = 1.$$

Goemans and Williamson (1995) proposed and analyzed the famous relaxation of this problem which consists in dropping the rank constraint, yielding a semidefinite program. Alternatively relaxing the rank constraint to be $\operatorname{rank}(X) \leq r$ for some $1 < r < n$ yields a tighter but nonconvex relaxation. Journée et al. (2010) observe that fixing the rank with the constraint $\operatorname{rank}(X) = r$ turns the search space into a smooth manifold, the fixed-rank elliptope, which can be optimized over using Riemannian optimization. In Manopt, simple code for this reads (with $Y \in \mathbb{R}^{n\times r}$ such that $X = YY^\top$):

```
% The problem structure hosts a manifold structure
% as well as function handles to define the cost
% function and its derivatives (here provided as
% Euclidean derivatives, which will be converted
% to their Riemannian equivalent).
problem.M = elliptopefactory(n, r);
problem.cost  = @(Y)     -trace(Y'*L*Y)/4;
problem.egrad = @(Y)     -(L*Y)/2;
problem.ehess = @(Y, U) -(L*U)/2;       % optional

% These diagnostics tools help make sure the
% gradient and Hessian are correct.
checkgradient(problem); pause;
checkhessian(problem);  pause;

% Minimize with trust-regions,
% a random initial guess and default options.
Y = trustregions(problem);
```

Randomly projecting Y as `s = sign(Y*randn(r, 1))` yields a cut. The Manopt distribution includes advanced code for this example, where the caching functionalities are used to avoid redundant computations of the product LY in the cost and the gradient, and the rank r is increased gradually to obtain a global solution of the max-cut SDP (and hence a formal upperbound), following the procedure in (Journée et al., 2010).

Acknowledgments

NB and BM are research fellows with the FNRS (aspirants). This paper presents research results of the Belgian Network DYSCO (Dynamical Systems, Control, and Optimization), funded by the Interuniversity Attraction Poles Programme initiated by the Belgian Science Policy Office. It was financially supported by the Belgian FRFC.

References

P.-A. Absil and K.A. Gallivan. Joint diagonalization on the oblique manifold for independent component analysis. In *Acoustics, Speech and Signal Processing, ICASSP 2006. IEEE International Conference on*, volume 5, 2006. doi: 10.1109/ICASSP.2006.1661433.

P.-A. Absil, C. G. Baker, and K. A. Gallivan. Trust-region methods on Riemannian manifolds. *Found. Comput. Math.*, 7(3):303–330, July 2007.

P.-A. Absil, R. Mahony, and R. Sepulchre. *Optimization Algorithms on Matrix Manifolds*. Princeton University Press, 2008.

S. Amari. Natural gradient learning for over- and under-complete bases in ICA. *Neural Computation*, 11(8):1875–1883, 1999.

L. Balzano, R. Nowak, and B. Recht. Online identification and tracking of subspaces from highly incomplete information. In *Communication, Control, and Computing (Allerton), 2010 48th Annual Allerton Conference on*, pages 704–711. IEEE, 2010.

S. Bonnabel. Stochastic gradient descent on Riemannian manifolds. *Automatic Control, IEEE Transactions on*, 58(9): 2217–2229, 2013. ISSN 0018-9286. doi: 10.1109/TAC.2013.2254619.

N. Boumal and P.-A. Absil. RTRMC: A Riemannian trust-region method for low-rank matrix completion. In J. Shawe-Taylor, R.S. Zemel, P. Bartlett, F.C.N. Pereira, and K.Q. Weinberger, editors, *Advances in Neural Information Processing Systems 24 (NIPS)*, pages 406–414. 2011.

N. Boumal, A. Singer, and P.-A. Absil. Robust estimation of rotations from relative measurements by maximum likelihood. *Proceedings of the 52nd Conference on Decision and Control, CDC 2013*, 2013.

N. Boumal, B. Mishra, P.-A. Absil, and R. Sepulchre. Manopt: a matlab toolbox for optimization on manifolds. *The Journal of Machine Learning Research*, 2014. Accepted for publication.

G. Dirr, U. Helmke, and C. Lageman. Nonsmooth Riemannian optimization with applications to sphere packing and grasping. In *Lagrangian and Hamiltonian Methods for Nonlinear Control 2006*, volume 366, pages 29–45. Springer, 2007.

A. Edelman, T.A. Arias, and S.T. Smith. The geometry of algorithms with orthogonality constraints. *SIAM journal on Matrix Analysis and Applications*, 20(2):303–353, 1998.

M.X. Goemans and D.P. Williamson. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *Journal of the ACM (JACM)*, 42(6):1115–1145, 1995.

I. Grubišić and R. Pietersz. Efficient rank reduction of correlation matrices. *Linear algebra and its applications*, 422(2): 629–653, 2007.

M. Journée, F. Bach, P.-A. Absil, and R. Sepulchre. Low-rank optimization on the cone of positive semidefinite matrices. *SIAM Journal on Optimization*, 20(5):2327–2351, 2010.

R.H. Keshavan, A. Montanari, and S. Oh. Matrix completion from noisy entries. *The Journal of Machine Learning Research*, 99:2057–2078, 2010.

D. Kressner, M. Steinlechner, and B. Vandereycken. Low-rank tensor completion by Riemannian optimization. Technical report, École polytechnique fédérale de Lausanne, 2013.

G. Meyer, S. Bonnabel, and R. Sepulchre. Regression on fixed-rank positive semidefinite matrices: a Riemannian approach. *The Journal of Machine Learning Research*, 12:593–625, 2011.

B. Mishra, G. Meyer, and R. Sepulchre. Low-rank optimization for distance matrix completion. In *Decision and Control and European Control Conference (CDC-ECC), 2011 50th IEEE Conference on*, pages 4455–4460. IEEE, 2011.

B. Mishra, G. Meyer, S. Bonnabel, and R. Sepulchre. Fixed-rank matrix factorizations and Riemannian low-rank optimization. *Arxiv preprint arXiv:1209.0430*, 2012.

W. Ring and B. Wirth. Optimization methods on Riemannian manifolds and their application to shape space. *SIAM Journal on Optimization*, 22(2):596–627, 2012.

U. Shalit, D. Weinshall, and G. Chechik. Online learning in the embedded manifold of low-rank matrices. *The Journal of Machine Learning Research*, 13:429–458, 2012.

F.J. Theis, T.P. Cason, and P.-A. Absil. Soft dimension reduction for ICA by joint diagonalization on the Stiefel manifold. In *Independent Component Analysis and Signal Separation*, pages 354–361. Springer, 2009.

R. Tron and R. Vidal. Distributed image-based 3D localization of camera sensor networks. In *Decision and Control, held jointly with the 28th Chinese Control Conference. Proceedings of the 48th IEEE Conference on*, pages 901–908. IEEE, 2009.

B. Vandereycken. Low-rank matrix completion by Riemannian optimization. *SIAM Journal on Optimization*, 23(2):1214–1236, 2013. doi: 10.1137/110845768.

Bayesian Estimation of Sparse Smooth Speckle Shape Models for Motion Tracking in Medical Ultrasound

Shaun Bundervoet, Colas Schretter, Ann Dooms and Peter Schelkens

Vrije Universiteit Brussel, Dept. of Electronics and Informatics, Pleinlaan 2, B-1050 Brussels, Belgium

iMinds, Dept. of Multimedia Technologies, Gaston Crommenlaan 8 box 102, B-9050 Ghent, Belgium

***Abstract*— Emerging ultrasound phased-array technologies will soon enable the acquisition of high-resolution 3D+T images for medical applications. Processing the huge amount of spatiotemporal measurements remains a practical challenge. In this work, dynamic ultrasound images are sparsely represented by a mixture of moving speckles. We model the shape of a speckle and its locally linear motion with a weighted multivariate Gaussian kernel. Parameters of the model are estimated with online Bayesian learning from a stream of random measurements. In our preliminary experiments with a simulated phantom of a moving cylindrical structure, the optical flow of speckles is estimated for a vertical line profile and compared to the ground truth. The mean accuracy of the linear motion estimate is of 93.53%, using only a statistically sufficient random subset of the data.**

1 Introduction

Measuring strain in structural soft tissue, such as tendons, could be an invaluable diagnostic tool for physicians [1]. For example, in orthopedics, measuring strain in ligaments will assist the placement procedure of joint implants and ultimately decrease the number of follow-up procedures [2]. Recent developments in ultrasound (US) imaging allows for the acquisition of high resolution volumetric images. This technology enables researchers to leverage existing imaging techniques such as elastography [3] or model-based biomechanical simulations of deformation [4] to the problem of assessing strain in tendons. Unfortunatelly, these techniques are not yet capable of capturing local motion occurring in small scale soft tissues.

In this work, we aim at adapting the speckle tracking echocardiography (STE) technique [5] from cardiology to the prescribed situation. Therefore, we propose a novel method for directly estimating motion of a sparse selection of speckle shape models instead of using detection and explicit tracking [6]. Speckles are forming the characteristic granular appearance in US images, which is due to deterministic interference originating from sub-wavelength scattering sites [7]. We approximate speckles and their local linear motion with a mixture of weighted multivariate Gaussian kernels. A Bayesian method is used for parametric estimation of the mixture components, which are representing the optical flow that is traced by the trajectories of speckles. Preliminary results are shown for simulated one-dimensional vertical line profiles.

2 Data and Image Models

We collect echography measurements in a dense discrete spatiotemporal amplitude field $A(x,t)$, with $x \in \{1,\dots,N_x\}$ and $t \in \{1,\dots,N_t\}$. This matrix of real positive amplitude values is first normalized for every time frame t such that

$$\sum_{i=1}^{N_x} A(x_i,t) = 1,\ \forall t \in \{1,\dots,N_t\}.$$

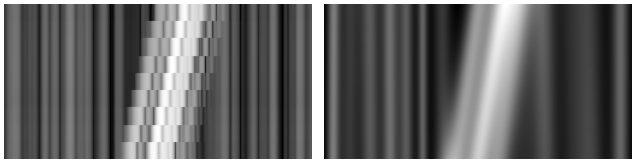

Spatio-temporal data Continuous reconstruction

Figure 1: Optical flow images of the input data consisting of 2597 echo measurements over 9 discrete time frames (left) and the smooth continuous spatiotemporal reconstruction using 16 multivariate mixture components (right).

We approximate $A(x,t)$ with a continuous and smooth conditional Gaussian mixture model $G(x,t)$ having the form

$$A(x,t) \approx G(x,t) = \sum_{k=1}^{K} w_k \, \frac{g(x,t;\Theta_k)}{g(t;\Theta_k)},$$

with the set of model parameters

$$\Theta = \{w_k, \Theta_k\}_{k=1}^{K} \quad \text{with} \quad \Theta_k = \left(\mu, \tau, \sigma_x^2, \sigma_t^2, \sigma_{xt}\right)_k$$

where $g(x,t;\Theta_k)$ can be decomposed as the product of a marginal (temporal) and the conditional (spatial) distributions:

$$g(x,t;\Theta_k) = g(t;\tau,\sigma_t^2) \times g(x;\bar{\mu},\bar{\sigma}_x^2),$$

where the parameters of the conditional distribution are

$$\bar{\mu} = \mu + \frac{\sigma_{xt}}{\sigma_t^2}(t-\tau) \ \text{ and } \ \bar{\sigma}_x^2 = \sigma_x^2 - \frac{\sigma_{xt}}{\sigma_t^2}\sigma_x^2,$$

and the Gaussian kernel $g(x;\mu,\sigma^2)$ is defined by

$$g(x;\mu,\sigma^2) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{1}{2}\frac{(x-\mu)^2}{\sigma^2}\right).$$

Before estimating the mixture parameters, we first sample $A(x,t)$ by taking a subset of $M \ll N_x \times N_t$ amplitude values

$$D = \{(x_m, t_m, a_m)\}_{m=1}^{M} \ \text{ with } \ a_m = A(\lfloor x_m \rfloor, \lfloor t_m \rfloor)$$

where (x_m, t_m) are drawn from the two-dimensional uniform random distribution $\mathcal{U}(1,N_x) \times \mathcal{U}(1,N_t)$.

3 Reconstruction

The set Θ of all weights and parameters of the mixture model $G(x,t)$ is now estimated by maximizing the log-likelihood

$$L(\Theta;D) = P(D;\Theta) = \sum_{m=1}^{M} \log \sum_{k=1}^{K} w_k \, \frac{g(x_m,t_m;\Theta_k)}{g(t_m;\Theta_k)}.$$

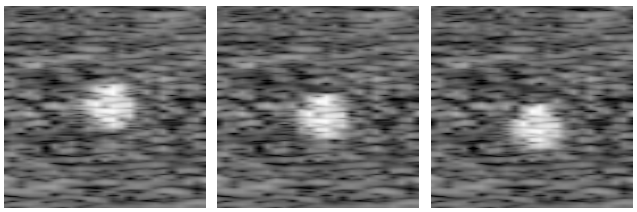

No displacement 1mm displacement 2mm displacement

Figure 2: Three selected frames of the moving cylinder phantom. The total course of the cylinder was 2mm long, such that the measured global displacement of the center of mass was 1.308mm.

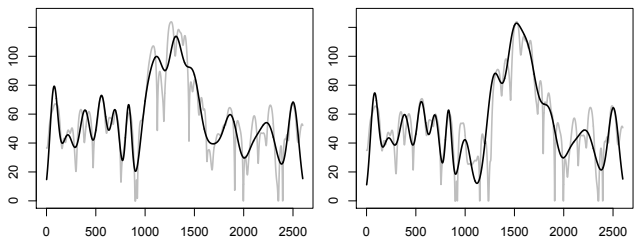

Figure 3: Plot of original (gray) and reconstructed (black) vertical line profiles for two selected time frames corresponding to the start (left) and end (right) of dynamic data acquisition. The sparse reconstruction uses 16 kernels.

We use the non-iterative online expectation-maximization (EM) method [8, 9] for estimating the parameters. Online EM alternates between performing an expectation (E-step) which evaluates membership probabilities for the log-likelihood using the current estimate of the parameters, and a maximization (M-step), which updates parameters for maximizing the expected log-likelihood found in the E-step. The model is first initialized with a regular grid, then the two E- and M- steps are performed successively for every sample, converging to the final model.

E-step When updating the model from the information carried by the m-th sample (x_m, t_m, a_m), we first compute the membership probability $p_k(x_m|t_m)$, expressing the probability that the sample is drawn from the conditional distribution of the component k at given time t_m using

$$p_k(x_m|t_m) = \frac{w_k \; g(x_m; \bar{\mu}, \bar{\sigma}_x^2)}{\sum_{k=1}^{K} w_k \; g(x_m; \bar{\mu}, \bar{\sigma}_x^2)},$$

where the parameters $\bar{\mu}$ and $\bar{\sigma}_x^2$ of the conditional distribution are calculated as previously described.

M-step In the second step, the parameters are updated such that the likelihood of observing the new sample is improved by applying the following update rules:

$$\begin{cases} w_k \leftarrow w_k \, (1+\alpha), \\ \mu \leftarrow \mu + \delta_x \, \alpha, \\ \tau \leftarrow \tau + \delta_t \, \alpha, \end{cases} \quad \begin{cases} \sigma_x^2 \leftarrow (\sigma_x^2 + \delta_x^2 \, \alpha) \, (1-\alpha), \\ \sigma_t^2 \leftarrow (\sigma_t^2 + \delta_t^2 \, \alpha) \, (1-\alpha), \\ \sigma_{xt} \leftarrow (\sigma_{xt} + \delta_x \, \delta_t \, \alpha) \, (1-\alpha), \end{cases}$$

for each $k \in \{1, \ldots, K\}$ with

$$\alpha = \frac{p_k(x_m|t_m) \, a_m}{w_k} \quad \text{and} \quad \begin{cases} \delta_x = x_m - \mu_k, \\ \delta_t = t_m - \tau_k. \end{cases}$$

Initialization We initialize each component of a given index k with the following weight $w_k = N_t/K$ and parameters

$$\mu = (k - \tfrac{1}{2}) \, \tfrac{N_x}{K}, \; \tau = \tfrac{N_t}{2}, \sigma_x^2 = \tfrac{N_x^2}{12K^2}, \; \sigma_t^2 = \tfrac{N_t^2}{12}, \; \sigma_{xt} = 0.$$

As presented, the weights are systematically incremented for every new sample in the M-step. In order to adapt smoothly to new data, a "forgetting" mechanism should be used to adjust parameters to the recent sufficient statistics. We simply apply a multiplicative decay factor $\gamma = 1 - 1/N_t$ to every component such that the total mixture weight remains normalized to N_t.

4 Experiment

For this paper we use the Field II program [7] in order to simulate RF-signals from a linear US probe consisting of 192 (64 active) elements. To this end a 3D phantom is generated containing a 10mm long high-intensity cylinder of 4mm radius. The whole numerical phantom consists of 5000 uniformly distributed and amplitude-weighted random scatterers for a region of $30 \times 30 \times 10\text{mm}^3$. This cylinder is then consecutively moved downwards by 8 consecutive steps of 0.25mm along the vertical axis, while the background is left unaltered.

For the simulation, the probe is positioned 15mm from the phantom emitting a 3.5Mhz pulse focused at a depth of 30mm. Afterwards the return signals is sampled at a frequency of 100Mhz. Using this method a total of 25 RF lines are created on which we perform log envelope detection. The images shown in Figure 2 show three temporal frames of the input data after downsampling along the axial dimension by a factor 10.

For the experiment we created a 2D optical flow image $A(x,t)$ using the central vertical line profile line at all $N_t = 9$ temporal steps which is shown in Figure 1. The number of amplitude measurements for the profile lines is $N_x = 2597$, the number of mixture components was set to $K = 16$, and the average number of samples used for statistical estimations of each component was set to was set to $N = 512$ such that the total number of samples was $M = NK = 8192$.

Reconstructed line profiles in Figure 3 show that the linear motion of speckles covering the cylinder is successfully captured with 16 mixture components. The center of mass in the data shifted by 1.308mm between the start and end of the simulation. The reconstructed displacement was of 1.224mm, reaching a relative error of 6.47% in this experiment. Similar experiments have been conducted with 8 and 32 mixture components. The motion extraction was not accurate with only 8 components, while we obtained comparable results with 32 components. We observed that the simulated speckle pattern slightly deforms over time while following the displacement path of the object. This effect has not yet been accounted for.

5 Conclusion

We represent dynamic ultrasound images as a mixture of moving speckle shape models that are approximated by smooth weighted multivariate Gaussian kernels. Maximum likelihood values of these kernel parameters can therefore be estimated by a conditional expectation-maximization method, using simple closed-form update rules. Experiments have been conducted on a line profile of a simple cylinder phantom. We observed that the sparse reconstruction approximates smoothly the acquired data. The subtle motion of the moving structure on a static background was successfully captured in the model. Future work will consider modeling the space-variant point spread function of speckles and extend the method to 3D+T imaging for assessing results on clinical data.

Acknowledgment

This work is supported by the iMinds ICON 3DUS project and the VUB Strategic Research Programme M^3D^2. S. Bundervoet is supported by the grant nr.131813 from the Agency for Innovation by Science and Technology in Flanders (IWT).

References

[1] T. J. Carter, M. Sermesant, D. M. Cash, D. C. Barratt, C. Tanner, D. J. Hawkes, "Application of soft tissue modelling to image-guided surgery", Med. Eng. Phys. vol. 27, no. 10, pp. 893–909, 2005.

[2] L. Scheys, P. Slagmolen, F. Morestin et al., "Non-invasive measurement of regional intratendinous strain using dynamic ultrasound: An ex vivo validation experiment in porcine patellar tendon", Proc. of ISTA, Bruges, 20-23 September, 2011.

[3] L. A. Chernak and D. G. Thelen, "Tendon motion and strain patterns evaluated with two-dimensional ultrasound elastography", Journal of Biomechanics, vol. 45, no. 15, pp. 2618–2623, 2012.

[4] T. S. Pheiffer, R. C. Thompson, D. Rucker, "Model-based correction of tissue compression for tracked ultrasound in soft tissue image-guided surgery", Ultrasound in Medicine & Biology, in press, 2014.

[5] J. DHooge, A. Heimdal, F. Jamal, T. Kukulski, B. Bijnens, F. Rademakers, et al., "Regional strain and strain rate measurements by cardiac ultrasound: principles, implementation and limitations", European Journal of Echocardiography, pp. 154–170, 2000.

[6] C. B. Compas, B. A. Lin, S. Sampath et al., "Multi-frame radial basis functions to combine shape and speckle tracking for cardiac deformation analysis in echocardiography", Functional Imaging and Modeling of the Heart, 6th International Conference, LNCS 6666, pp. 113–120, 2011.

[7] J. A. Jensen, "Field: a program for simulating ultrasound systems", Medical & Biological Engineering & Computing, vol. 34, pp. 351–353, 1996.

[8] D. Mike Titterington, "Recursive parameter estimation using incomplete data", Journal of the Royal Statistical Society. Series B (Methodological), vol. 46, no. 2, pp. 257–267, 1984.

[9] R. Neal and G. E. Hinton, "A view of the EM algorithm that justifies incremental, sparse, and other variants", Learning in Graphical Models, pp. 355–368, 1998.

Learning a fast transform with a dictionary

Olivier Chabiron[1], François Malgouyres[2], Jean-Yves Tourneret[1] and Nicolas Dobigeon[1].

[1]Institut de Recherche en Informatique de Toulouse, IRIT-CNRS UMR 5505, ENSEEIHT, Toulouse, France.

[2]Institut de Mathématiques de Toulouse, IMT-CNRS UMR 5219, Université de Toulouse, Toulouse, France.

***Abstract*— A powerful approach to sparse representation, dictionary learning consists in finding a redundant frame in which the representation of a particular class of images is sparse. In practice, all algorithms performing dictionary learning iteratively estimate the dictionary and a sparse representation of the images using this dictionary. However, the numerical complexity of dictionary learning restricts its use to atoms with a small support.**

A way to alleviate these issues is introduced in this paper, consisting in dictionary atoms obtained by translating the composition of K convolutions with S-sparse kernels of known support. The dictionary update step associated with this strategy is a non-convex optimization problem, which we study here.

A block-coordinate descent or Gauss-Seidel algorithm is proposed to solve this problem, whose search space is of dimension KS, which is much smaller than the size of the image. Moreover, the complexity of the algorithm is linear with respect to the size of the image, allowing larger atoms to be learned (as opposed to small patches). An experiment is presented that shows the approximation of a large cosine atom with $K = 7$ sparse kernels, demonstrating a very good accuracy.

1 Introduction

The problem we introduce in this paper is motivated by the dictionary learning (DL) field. DL was pioneered by [8, 9] and has received a growing attention since then. The principle behind DL is to find a representation for data that makes it simpler, sparser. We invite the reader to consult [4] for more details about sparse representations and DL. The archetype of the DL strategy is to look for a dictionary as the solution of the following optimization problem

$$\operatorname{argmin}_{\mathbf{D},(\mathbf{x}_i)_{1\leq i\leq I}} \quad \sum_{i=1}^{I} \|\mathbf{D}\mathbf{x}_i - \mathbf{y}_i\|_2^2 + f(\mathbf{x}_i),$$

where $\mathbf{y}_i$ are the learning database, $\mathbf{D}$ is the dictionary matrix, whose columns are the atoms, and f is a sparsity-inducing function. The resulting problem can be solved (or approximatively solved) by many methods including MOD [5] and K-SVD [1]. All these approaches rely on alternatively updating the codes $\mathbf{x}_i$ and the dictionary $\mathbf{D}$.

Our primary motivation for considering the observation model (1) comes from computational issues. Usually, DL is applied to small patches, because of the computational cost of repeatedly computing of the matrix-vector products $\mathbf{D}\mathbf{x}_i$, which is worse than $\mathcal{O}(N^2)$. Moreover, the cost of the dictionary update is usually worse than $\mathcal{O}(N^3)$.

We propose a model where the learned atoms are a composition of K convolutions with S-sparse kernels. The interest for such a constraint is to provide numerically effective dictionaries and allow to consider larger atoms. Indeed, the search space is only of dimension KS, which is typically smaller than the size of the target atom.

The present work focuses on the dictionary update step of one atom. In this context, the code $\mathbf{x}$ is known. Our goals are both to approximate a large target atom $\boldsymbol{\kappa}$ with our model and to obtain target atoms whose manipulation is numerically efficient. This translates into a non-convex optimization problem.

2 Problem formulation

Let consider an observed d-dimensional signal $\mathbf{y}$ of $(\mathbb{R}^N)^d$, assumed to result from the convolution of a known input signal $\mathbf{x} \in (\mathbb{R}^N)^d$ with an unknown target kernel $\boldsymbol{\kappa} \in (\mathbb{R}^N)^d$, contaminated by an additive noise $\mathbf{b}$ following the linear model

$$\mathbf{y} = \boldsymbol{\kappa} * \mathbf{x} + \mathbf{b}, \tag{1}$$

where $*$ stands for the circular discrete convolution[1] in dimension d. For instance, the unknown target kernel $\boldsymbol{\kappa} \in (\mathbb{R}^N)^d$ may refer to the unknown impulse response of a 1D-filter or, conversely, to the point spread function of a 2D-filtering operator.

The problem addressed in this paper consists of approximating the unknown kernel $\boldsymbol{\kappa}$ by a composition of convolutions with $K \geq 2$ sparse kernels $(\mathbf{h}^k)_{1\leq k\leq K} \in ((\mathbb{R}^N)^d)^K$

$$\boldsymbol{\kappa} \approx \hat{\boldsymbol{\kappa}} \triangleq \mathbf{h}^1 * \cdots * \mathbf{h}^K. \tag{2}$$

The kernels $\mathbf{h}^1, \ldots, \mathbf{h}^K$ are constrained to have less than a fixed number S of non-zero elements, i.e., they are assumed to be at most S-sparse. As stated before, this assumption aims at providing a cost-effective dictionary by reducing the computations for $\mathbf{x} * \boldsymbol{\kappa}$. Furthermore, the locations of their non-zero elements in $\{0, \ldots, N\}^d$ are assumed to be known or pre-set. More precisely, the support of the kth kernel (i.e., the locations of the non zero elements of $\mathbf{h}^k$), denoted $\operatorname{supp}(\mathbf{h}^k)$, is constrained to a fixed set of discrete indexes $\mathcal{S}_k$

$$\operatorname{supp}(\mathbf{h}^k) \subset \mathcal{S}_k \ , \ \forall k \in \{1, \ldots, K\} \tag{3}$$

An example of indexes for 1D convolution kernel would be

$$\mathcal{S}_k = \{k-1, 2k-1, \ldots, Sk-1\}. \tag{4}$$

Assuming that the noise vector $\mathbf{b}$ is an independent and identically distributed Gaussian sequence, approximating the unknown convolution kernel $\boldsymbol{\kappa}$ from the observed measurements $\mathbf{y}$ can be formulated as the following optimization problem

$$(P_0): \begin{cases} \operatorname{argmin}_{\mathbf{h}\in((\mathbb{R}^N)^d)^K} \|\mathbf{y} - \mathbf{h}^1 * \cdots * \mathbf{h}^K * \mathbf{x}\|_2^2, \\ \text{subject to } \operatorname{supp}(\mathbf{h}^k) \subset \mathcal{S}_k \ , \forall k \in \{1, \ldots, K\} \end{cases}$$

where $\|\cdot\|_2$ stands for the usual Euclidean norm in $(\mathbb{R}^N)^d$. The problem (P_0) is non convex. Thus, depending on the values of K, $(\mathcal{S}_k)_{1\leq k\leq K}$, $\mathbf{x}$ and $\mathbf{y}$, it might be difficult or impossible

[1]All the elements of $(\mathbb{R}^N)^d$ are extended over $\mathbb{Z}^d$ by periodization.

to find a good approximation of a global minimizer of (P_0). Our objectives are to study if such a problem bends itself to global optimization, and to assess the approximation power of the computed compositions of convolutions.

3 Block-coordinate descent

The problem (P_0), as formulated in the previous section, is unhandy for global optimization. As detailed in [2], it has irrelevant stationary points and is non-convex (though infinitely differentiable). To adress these issues, a scalar weight λ is introduced and kernels are constrained to have a unit norm. Moreover, we elect a block-coordinate formulation in order to solve the problem with a Gauss-Seidel type algorithm (called Alternate Least Squares, sharing many similarities with the one used in [7]).

$$(P_k): \left\{ \begin{array}{l} \operatorname{argmin}_{\lambda\in\mathbb{R},\mathbf{h}\in\mathbb{R}^N} \|\mathbf{y}-\lambda\mathbf{h}*\mathbf{x}^k\|_2^2, \\ \text{subject to } \operatorname{supp}(\mathbf{h})\subset\mathcal{S}_k \text{ and } \|\mathbf{h}\|_2=1 \end{array} \right.$$

with

$$\mathbf{x}^k = \mathbf{h}^1 * \cdots * \mathbf{h}^{k-1} * \mathbf{h}^{k+1} * \cdots * \mathbf{h}^K * \mathbf{x}, \qquad (5)$$

where the kernels $\mathbf{h}^{k'}$ are fixed $\forall k' \neq k$. The problem (P_k) is linear and can be expressed as a matrix-vector product considering only the elements of $\mathbf{h}$ that belong to its support: The idea is to alternatively solve (P_k) by iterating on k. The support constraint significantly reduces the search space of the problem, and thus the amount of calculations needed to solve it for a stationary point. Algorithm 1 shows an overview of the resolution of (P_k). The computational complexity associated with a passage in the while loop is $O((K+S)KSN^d)$, i.e., it is linear with respect to the size N^d of the signal. The detailed steps to solving (P_k) are given in [2].

Algorithm 1: ALS algorithm

Input:
$\mathbf{y}$: target measurements;
$\mathbf{x}$: known coefficients;
$(\mathcal{S}_k)_{1\leq k\leq K}$: supports of the kernels $(\mathbf{h}^k)_{1\leq k\leq K}$.
Output:
$(\mathbf{h}^k)_{1\leq k\leq K}$: convolution kernels such that
$\mathbf{h}^1 * \ldots * \mathbf{h}^K \approx \boldsymbol{\kappa}$.
begin
  Initialize the kernels $((\mathbf{h}_p^k)_{p\in N})_{1\leq k\leq K}$;
  while *not converged* **do**
    for *k = 1 ,... , K* **do**
      Update $\mathbf{h}^k$ and λ with a minimizer of (P_k) ;

4 Synthetic example

In this section, we show an experiment consisting of approximating a 2D cosine atom $\boldsymbol{\kappa}$ in an image $\mathbf{y}$ of size 64×64 (i.e., $d=2$ and $N=64$). Such an atom can be seen as a large local cosine or a Fourier atom, both widely used in image processing. The interest of this atom is that it covers the whole image and is of a rather large support, making it difficult to handle with existing dictionary learning strategies.

$$\boldsymbol{\kappa}_p = \cos\left(2\pi\frac{\langle p,(2,5)\rangle}{N}\right) \quad ,\forall p\in\{0,\ldots,63\}^2.$$

The code $\mathbf{x}$ is a sparse image whose elements are chosen independent and identically distributed according to a Bernoulli-Gaussian distribution, widely used in sparse signal and image deconvolution [3, 6, 10]. Therefore, $\mathbf{y}$ contains a few weighted translations of the cosine atom $\boldsymbol{\kappa}$[2]. The target $\mathbf{y}$ is built with additive Gaussian noise of variance $\sigma^2=0.5$. Kernel supports have been set to a simple 5×5 square, linearly dilated with k, similar to the 1-D example given in (4).

Figures 1 and 2 show the cosine image $\mathbf{y}$, its approximation $\lambda\mathbf{x}*\mathbf{h}^1*\cdots*\mathbf{h}^K$, the actual atom $\boldsymbol{\kappa}$ and $\lambda\mathbf{h}^1*\cdots*\mathbf{h}^K$, for $K=7$ and $S=25$. The results obtained here are quite accurate even though the cosine image was corrupted by additive noise.

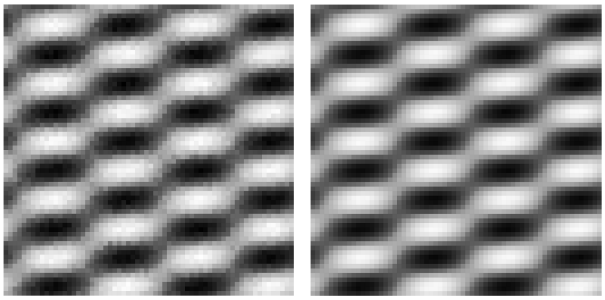

Figure 1: Cosine approximation with $K=7$, $S=25$, and Gaussian noise of variance $\sigma^2=0.5$. Cosine image $\mathbf{y}$ (left) and approximation $\lambda\mathbf{x}*\mathbf{h}^1*\cdots*\mathbf{h}^K$ (right).

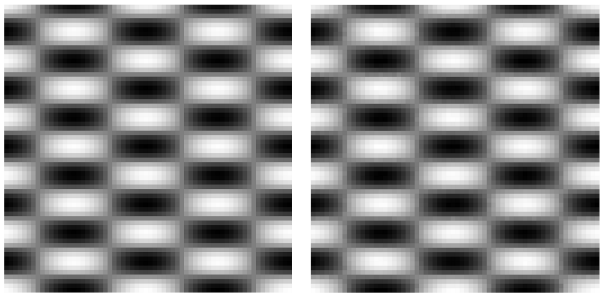

Figure 2: Cosine approximation with $K=7$, $S=25$, and Gaussian noise of variance $\sigma^2=0.5$. True atom $\boldsymbol{\kappa}$ (left) and approximation $\lambda\mathbf{h}^1*\cdots*\mathbf{h}^K$ (right).

5 Conclusion

This work shows that simple atoms can be accurately approximated with a composition of convolutions. The kernels used in the approximation are constrained to be sparse (i.e., with sparse supports), leading to a computationally efficient algorithm, despite the non-convexity of the function to optimize. This efficiency was illustrated on a 2D-cosine function, but similar experiments conducted with archetypal kernels (e.g., wavelets or curvelets) show similar performances [2].

The proposed modeling and algorithmic schemes open new perspectives on the general problem of dictionary learning. More specifically, it seems reasonable to derive a DL technique which recovers large structured dictionary whose atoms consist of compositions of convolutions.

Finally, how to choose, set or draw the kernel supports remains a large and yet unexplored issue, that may have significant impact on the method performances.

[2] A sum of cosines of same frequency and different phases will yield a cosine of unchanged frequency.

References

[1] M. Aharon, M. Elad, and A. M. Bruckstein. The K-SVD, an algorithm for designing overcomplete dictionaries for sparse representation. *IEEE Trans. Signal Process.*, 54(11):4311–4322, 2006.

[2] Olivier Chabiron, Francois Malgouyres, Jean-Yves Tourneret, and Nicolas Dobigeon. Toward fast transform learning. 2013. submitted to IJCV.

[3] F. Champagnat, Y. Goussard, and J. Idier. Unsupervised deconvolution of sparse spike trains using stochastic approximation. *IEEE Trans. Signal Process.*, 44(12):29882998, 1996.

[4] M. Elad. *Sparse and redundant representations: From theory to applications in signal and image processing*. Springer, 2010.

[5] K. Engan, S. O. Aase, and J. Hakon Husoy. Method of optimal directions for frame design. In *Proc. IEEE Int. Conf. Acoust., Speech, and Signal Processing (ICASSP)*, pages 2443–2446, Washington, DC, USA, 1999.

[6] G. Kail, J.-Y. Tourneret, N. Dobigeon, and F. Hlawatsch. Blind deconvolution of sparse pulse sequences under a minimum distance constraint: A partially collapsed Gibbs sampler method. *IEEE Trans. Signal Process.*, 60(6):2727–2743, june 2012.

[7] L. De Lathauwer, B. De Moor, and J. Vandewalle. On the best rank-1 and rank-(r1,r2,. . .,rn) approximation of higher-order tensors. *SIAM J. Matrix Anal. Appl.*, 21(4):1324–1342, 2000.

[8] M. S. Lewicki and T. J. Sejnowski. Learning overcomplete representations. *Neural Computation*, 12(2):337–365, 2000.

[9] B. A. Olshausen and D. J. Field. Sparse coding with an overcomplete basis set: A strategy employed by v1? *Vision Research*, 37(23):3311 – 3325, 1997.

[10] C. Quinsac, N. Dobigeon, A. Basarab, J.-Y. Tourneret, and D. Kouamé. Bayesian compressed sensing in ultrasound imaging. In *Proc. of Third International Workshop on Computational Advances in Multi-Sensor Adaptive Processing (CAMSAP11)*, San Juan, Puerto Rico, 2011.

A diffusion strategy for distributed dictionary learning

Pierre Chainais[1], Cédric Richard[2].

[1]LAGIS UMR CNRS 8219 - Ecole Centrale Lille, INRIA-Lille Nord Europe, SequeL, France
pierre.chainais@ec-lille.fr.

[2]Laboratoire Lagrange UMR CNRS 7293, University of Nice Sophia-Antipolis, France
cedric.richard@unice.fr

***Abstract*— We consider the problem of a set of nodes which is required to collectively learn a common dictionary from noisy measurements. This distributed dictionary learning approach may be useful in several contexts including sensor networks. Diffusion cooperation schemes have been proposed to estimate a consensus solution to distributed linear regression. This work proposes a diffusion-based adaptive dictionary learning strategy. Each node receives measurements which may be shared or not with its neighbors. All nodes cooperate with their neighbors by sharing their local dictionary to estimate a common representation. In a diffusion approach, the resulting algorithm corresponds to a distributed alternate optimization. Beyond dictionary learning, this strategy could be adapted to many matrix factorization problems in various settings. We illustrate its efficiency on some numerical experiments, including the difficult problem of blind hyperspectral images unmixing.**

1 Introduction

In a variety of contexts, huge amounts of high dimensional data are recorded from multiple sensors. When sensor networks are considered, it is desirable that computations be distributed over the network rather than centralized in some fusion unit. Indeed, centralizing all measurements lacks robustness - a failure of the central node is fatal - and scalability due to the needed energy and communication resources. In distributed computing, every node communicates with its neighbors only and processing is carried out by every node in the network. Another important remark is that relevant information from the data usually lives in a space of much reduced dimension compared to the physical space. The extraction of this relevant information calls for the identification of some adapted sparse representation of the data. Learning an adaptive *sparse representation* of the data using a *redundant dictionary* is useful for many tasks such as storing, transmitting or analyzing the data to understand its content, see [1] for an up-to-date review. Furthermore, the problem of dictionary learning belongs to the more general family of matrix factorization problems that appears in a host of applications. We study the problem of dictionary learning distributed over a sensor network in a setting where a set of nodes is required to collectively learn an adaptive sparse representation from independent observations. We consider the situation where a set of connected nodes records data from observations of the same kind of physical system: each observation is assumed to be described by a sparse representation using a common dictionary over all sensors. For instance, a set of cameras observe the same kind of scenes or a set of microphones records the same kind of sound environment.

The dictionary learning and the matrix factorization problems are connected to the linear regression problem. Indeed, the classical approach based on alternate minimization on the coefficients $\mathbf{X}$ and the dictionary $\mathbf{D}$ solves two linear regression problems knowing respectively $\mathbf{D}$ or $\mathbf{X}$. Several recent works have proposed efficient solutions to the problem of least mean square (LMS) distributed linear regression, see [2] and references therein. The main idea is to use a so-called *diffusion strategy*: each node n carries out its own estimation $\mathbf{D}_n$ of the same underlying linear regression vector $\mathbf{D}$ but can communicate with its neighbors as well. The information provided to some node by its neighbors is taken into account according to weights interpreted as diffusion coefficients. Under some mild conditions, the performance of such an approach in terms of mean squared error is similar to that of a centralized approach [3]. Let $\mathbf{D}_c$ the centralized estimate which uses all the observations at once. It can be shown that the error $\mathbf{E}\|\mathbf{D}_n - \mathbf{D}\|_2$ of the distributed estimate is of the same order as $\mathbf{E}\|\mathbf{D}_c - \mathbf{D}\|_2$: diffusion networks match the performance of the centralized solution.

Our work [4] gives strong indication that the classical dictionary learning technique based on block coordinate descent on the dictionary $\mathbf{D}$ and the coefficients $\mathbf{X}$ can be adapted to the distributed framework by adapting the diffusion strategy mentionned above. Our numerical experiments also strongly support this idea. Note that solving this type of matrix factorization problems is really at stake since it corresponds to many inverse problems: denoising, adaptive compression, recommendation systems... A distributed approach is highly desirable both for use in sensor networks and for parallelization of numerically expensive learning algorithms. In a second step, we may consider the more general situation where observations may also be shared between connected nodes.

2 Problem formulation

Many nodes, one dictionary. Consider N nodes over some region. In the following, boldfaced letters denote column vectors, and capital letters denote matrices. The node n takes q_n measurements $\mathbf{y}_n(i)$, $1 \leq i \leq q_n$ from some physical system. All the observations are assumed to originate from independent realizations $\mathbf{s}_n(i)$ of the same underlying stochastic source process $\mathbf{s}$. Each measurement is a noisy measurement

$$\mathbf{y}_n(i) = \mathbf{s}_n(i) + \mathbf{z}_n(i) \tag{1}$$

where $\mathbf{z}$ denotes the usual i.i.d. Gaussian noise with covariance matrix $\Sigma_n = \sigma_n^2 \mathbb{I}$. Our purpose is to learn a common redundant dictionary $\mathbf{D}$ which carries the characteristic properties of the data. This dictionary must yield a sparse representation of $\mathbf{s}$ so that:

$$\forall n, \quad \mathbf{y}_n(i) = \underbrace{\mathbf{D}\mathbf{x}_n(i)}_{\mathbf{s}_n(i)} + \mathbf{z}_n(i) \tag{2}$$

where $\mathbf{x}_n(i)$ features the coefficients $x_{nk}(i)$ associated to the contribution of atom $\mathbf{d}_k$, the k-th column in the dictionary ma-

trix $\mathbf{D}$, to $\mathbf{s}_n(i)$. The sparsity of $\mathbf{x}_n(i)$ means that only few components of $\mathbf{x}_n(i)$ are non zero.

We consider the situation where a unique dictionary $\mathbf{D}$ generates the observations at all nodes. On the contrary, observations will first not be shared between nodes (this is one potential generalization). Our purpose is to learn this dictionary in a distributed manner thanks to in-network computing only. As a consequence, each node will locally estimate a local dictionary $\mathbf{D}_n$ thanks to i) its observations $\mathbf{y}_n$ and ii) communication with its neighbors. The neighborhood of node n will be denoted by $\mathcal{N}_n$, including node n itself.

Dictionary learning. Various approaches to dictionary learning have been proposed [1]. Usually, in the centralized setting, the q observations are denoted by $\mathbf{y}(i) \in \mathbb{R}^p$ and grouped in a matrix $\mathbf{Y} = [\mathbf{y}(1), ..., \mathbf{y}(q)]$. As a consequence, $\mathbf{Y} \in \mathbb{R}^{p \times q}$. The dictionary (associated to some linear transform) is denoted by $\mathbf{D} \in \mathbb{R}^{p \times K}$: each column is one atom $\mathbf{d}_k$ of the dictionary. The coefficients associated to observations are $\mathbf{X} = [\mathbf{x}(1), ..., \mathbf{x}(q)]$. We will consider learning methods based on block coordinate descent or alternate optimization on $\mathbf{D}$ and $\mathbf{X}$ with a sparsity constraint on $\mathbf{X}$ [5, 1, 6]. The data is represented as the sum of a linear combination of atoms and a noise term $\mathbf{Z} \in \mathbb{R}^{p \times q}$:

$$\mathbf{Y} = \mathbf{D}\mathbf{X} + \mathbf{Z} \tag{3}$$

In the most usual setting featuring white Gaussian noise, one wants to solve:

$$(\mathbf{D}, \mathbf{X}) = \text{argmin}_{(\mathbf{D},\mathbf{X})} \quad \frac{1}{2}||\mathbf{Y} - \mathbf{D}\mathbf{X}||_2^2 + \lambda||\mathbf{X}||_1 \tag{4}$$

Under some mild conditions, this problem is known to provide a solution to the underlying L0-penalized problem [7].

3 Distributed alternate optimization for dictionary learning

Algorithm. The Adapt Then Combine diffusion strategy [2] for distributed estimation originates the following approach to distributed alternate optimization for dictionary learning. Diffusion is ensured by the communication between nodes sharing their dictionary estimate with neighbors in $\mathcal{N}_n$. Observations are taken simultaneously at each node so that a whole data matrix $\mathbf{Y}_n$ is assumed to be available at node n. Here index i stands for iterations. The case where data arrive sequentially at each node can also be dealt with at the price of a natural adaptation of the present approach. Each node must estimate both its local dictionary $\mathbf{D}_n$ and the coefficients $\mathbf{X}_n$ which describe observations $\mathbf{Y}_n = \mathbf{D}_n\mathbf{X}_n + \mathbf{Z}_n$. At each iteration i, only the local dictionary estimates $\mathbf{D}_{n,i}$ are assumed to be shared between neighbors, not observations. In summary, sparse representations are computed locally. Then each node updates its dictionary as a function of its local observations $\mathbf{Y}_n$ (Adapt step) and its neighbors' dictionaries (Combine step). Based on known results for the ATC strategy in its usual setting, we expect that Algorithm 1 below converges to an accurate estimate of the common underlying dictionary $\mathbf{D}$. Various choices can be considered for $\mathbf{A}$ such as some a priori fixed matrix $\mathbf{A}$ or with the relative degree variance (ν_ℓ = degree of node ℓ):

$$a_{\ell,n} = \frac{\nu_\ell \sigma_\ell^2}{\sum_{m \in \mathcal{N}_n} \nu_m \sigma_m^2} \tag{5}$$

Initialize $\mathbf{D}_{n,0}$, $\forall n$ (random subset of K observations $y_n(i)$).
Given a matrix $\mathbf{A}$ satisfying $\mathbf{1}^T\mathbf{A} = \mathbf{1}^T$, $i = 0$,

Repeat until convergence of $(\mathbf{D}_{n,i}, \mathbf{X}_{n,i})_{n=1:N}$

For each node n **repeat:**

1) Optimization w.r.t. $\mathbf{X}_{n,i}$ (sparse coding):
Given the dictionary $\mathbf{D}_{n,i}$, the coefficients $\mathbf{X}_{n,i}$ are estimated using a sparse coding method (Basis Pursuit, OMP, FOCUSS,...)

2) Optimization w.r.t. $\mathbf{D}_{n,i}$ (dictionary) e.g. by gradient descent:

$$\begin{cases} \boldsymbol{\psi}_{n,i+1} &= \mathbf{D}_{n,i} + \mu_n^D(\mathbf{Y}_n - \mathbf{D}_{n,i}\mathbf{X}_{n,i})\mathbf{X}_{n,i}^T \\ \mathbf{D}_{n,i+1} &= \sum_{\ell \in \mathcal{N}_k} a_{\ell,n}^D \boldsymbol{\psi}_{\ell,i} \text{ (diffusion)} \end{cases}$$

and $\forall 1 \leq k \leq K, \mathbf{d}_k \leftarrow \frac{\mathbf{d}_k}{\|\mathbf{d}_k\|_2}$ (normalization)

$\boldsymbol{\psi}_{n,i+1}$ can also be updated by MOD or K-SVD,...at node n.

EndFor (n)
$i \leftarrow i + 1$
EndRepeat

Numerical experiments. We present some numerical experiments to illustrate the relevance and efficiency of our approach. For instance we show the results obtained on a dataset built from a redundant random dictionary of 48 atoms of dimension 16 corresponding to image patches of size 4×4. Each data $y_n(i)$ is the linear combination of 3 atoms with i.i.d. coefficients uniformly distributed over $[-0.5, 0.5]$; various Gaussian noise levels are considered. We show that a set of 4 nodes in a symmetrically connected network consistently learn the same dictionary of 4×4 patches with good accuracy (45 atoms of the initial dictionary of 48 atoms are recovered with $\langle \mathbf{d}_j, \mathbf{d}_j^{(o)} \rangle \geq 0.99$).

We will also show the results of an application to the problem of blind unmixing of hyperspectral images. In this application, the network is simply made of connections between pixels which are spatially close or which carry similar spectral information. The graph underlying (hyper-)pixels makes them collaborate to learn spectral endmembers.

4 Conclusion

We present an original algorithm which solves the problem of distributed dictionary learning over a sensor network. This is made possible thanks to a diffusion strategy which permits local communication between neighbors. Connected nodes exchange their local dictionaries estimated from disjoint subsets of data. This algorithm adapts usual dictionary learning techniques for sparse representation to the context of in-network computing. This approach to the general problem of distributed matrix factorization paves the way towards many prospects and applications. Moreover, as far as computational complexity is concerned, distributed parallel implementations are a potentially interesting alternative to online learning techniques [8]. We may even consider a dynamical context where observations arrive over time so that the dictionary would also be learnt dynamically.

References

[1] I. Tosic and P. Frossard, "Dictionary learning," *IEEE Signal Processing Magazine*, vol. 28, no. 2, pp. 27 –38, march 2011.

[2] F. Cattivelli and A. Sayed, "Diffusion lms strategies for distributed estimation," *Signal Processing, IEEE Transactions on*, vol. 58, no. 3, pp. 1035 –1048, march 2010.

[3] X. Zhao and A. H. Sayed, "Performance limits for distributed estimation over lms adaptive networks," *IEEE Transactions on Signal Processing*, vol. 60, no. 10, pp. 5107–5124, 2012.

[4] P. Chainais and C. Richard, "Learning a common dictionary over a sensor network," in *Proc. of IEEE CAMSAP 2013*, 2013.

[5] P. Tseng, "Convergence of a block coordinate descent method for nondifferentiable minimization," *Journal of Optimization Theory and Applications*, vol. 109, pp. 475–494, 2001. [Online]. Available: http://dx.doi.org/10.1023/A%3A1017501703105

[6] M. Aharon, M. Elad, and A. Bruckstein, "K-SVD: An algorithm for designing overcomplete dictionaries for sparse representation," *IEEE Transactions on Signal Processing*, vol. 54, no. 11, pp. 4311 –4322, nov. 2006.

[7] J.-L. Starck, F. Murtagh, and J. Fadili, *Sparse Image and Signal Processing: Wavelets, Curvelets, Morphological Diversity*. Cambridge University Press, 2010.

[8] J. Mairal, F. Bach, J. Ponce, and G. Sapiro, "Online learning for matrix factorization and sparse coding," *J. Mach. Learn. Res.*, vol. 11, pp. 19–60, Mar. 2010. [Online]. Available: http://dl.acm.org/citation.cfm?id=1756006.1756008

Bayesian crack detection in high resolution data

Bruno Cornelis[1,2], Ann Dooms[1,2], Ingrid Daubechies[3] and David Dunson[4].

[1]ETRO (Department of Electronics and Informatics), Vrije Universiteit Brussel, Belgium.

[2]iMinds VZW, Gaston Crommenlaan 8 box 102, 9050 Ghent, Belgium.

[3]Department of Mathematics, Duke University, Durham, North Carolina, USA.

[4]Department of Statistical Science, Duke University, Durham, North Carolina, USA.

***Abstract*— We propose a semi-supervised crack detection method that can be used for high-dimensional and multimodal acquisitions of paintings. Our dataset consists of a recent collection of images of the Ghent Altarpiece (1432), one of Northern Europe's most important art masterpieces. We build a classifier that is able to discern crack pixels from the background consisting of non-crack pixels, making optimal use of the information that is provided by each modality. To accomplish this we employ a recently developed non-parametric Bayesian classifier, that uses tensor factorizations to characterize any conditional probability. A prior is placed on the parameters of the factorization such that every possible interaction between predictors is allowed while still identifying a sparse subset among these predictors.**

1 Introduction

The digitization of cultural artifacts is becoming an increasingly popular practice. Consequently, their digital analysis emerged as a key *non-invasive* tool for art historians and preservation experts. In this work, we focus on crack detection in the *Ghent Altarpiece* (1432). Being able to accurately detect cracks is of major importance in the context of art conservation since cracking is one of the most common forms of deterioration found in paintings. For most 15$^{\text{th}}$ century Flemish paintings on Baltic oak, humidity fluctuations cause the wooden support to shrink or expand and are the main cause for crack formation.

Previous efforts on crack detection in the Ghent Altarpiece can be found in [1] and an extension in [2], where three detection methods (filtering with elongated filters, a multi scale top hat transform and K-SVD for crack enhancement) are proposed. The images that we analyzed back then are scans of original photographic negatives. However, thanks to the project *Lasting Support*[1] (2010-2011), a new *multimodal* dataset of ultra high resolution was made available. The panels of the altarpiece were documented with digital macrophotography in both the visible and infrared parts of the electromagnetic spectrum, together with previously acquired X-rays.

This multimodal dataset requires new crack detection techniques that are able to combine the information provided by each modality. Simply applying the methods from [2], which were designed to work well on one single modality, would require choosing an additional set of parameters per modality. The goal is to build an accurate classifier and select a subset of most important predictors from a a high dimensional collection of categorical predictor vectors. We define a model, by using a carefully-structured *Tucker factorization*, that can characterize any conditional probability and facilitates variable selection and the modeling of higher-order interactions [3]. We follow a Bayesian approach and use a Markov chain Monte Carlo (MCMC) algorithm for posterior computation and hence allow for uncertainties in the to be included predictors. The proposed approach which we refer to as *Bayesian conditional tensor factorization* or *BCTF*, introduced in [4], shows comparable detection results to the Random Forest (RF) algorithm [5] but has the tendency to produce finer crack maps and is faster, which is an attractive property when working with ultra high-dimensional data.

[1] `http://closertovaneyck.kikirpa.be`

2 Data preprocessing

The available modalities of the painting are fundamentally very different and therefore direct registration is a challenging task. Since cracks are the only components that are more or less consistent throughout the modalities, crude crack maps are initially constructed. These are obtained by first filtering the image with elongated filters [2] and a subsequent thresholding step produces binary *crack maps* that show an initial estimate of the crack locations. These are registered using the algorithm described in [6] and the resulting transformation is then applied to the original images.

The penetrating nature of X-rays make them an appropriate modality for the detection of cracks, since they usually appear clearly delineated. However, the presence of the wood grain in some of the panels can act as a confounding factor during the detection process. We use a texture separation technique, called morphological component analysis (MCA) to separate the texture from the natural part in images. The textural part in our images consists of the wood grain, for which we select a (Local) Discrete Cosine Transform (DCT) dictionary. The Dual Tree Complex Wavelet dictionary is used for the representation of piecewise smooth content.

3 Image features

Each modality is processed and filtered with a wide range of filters commonly used in image processing. The methods and filters that are used to construct the image predictor vectors include elongated filters, a Frangi Vesselness filter [7], the eigenvalues of a structure tensor, black Top Hat transforms with different structuring elements [8], Local Binary Patterns [9], pixel intensities, median filters, Laplacian of Gaussian (LoG) filters and the *Leung-Malik* filter bank, a multi-scale, multi-orientation filter bank of 48 filters. Each pixel is now characterized by a feature vector of length 208 that is the result of all the operations described above, applied to each modality independently. Once these are extracted they are quantized into an experimentally chosen number of bins d and used in the subsequent classification step, described below.

4 Bayesian Conditional Tensor Factorizations

We use a recently developed framework for nonparametric Bayes classification through tensor factorizations of the conditional probability $P(Y|X)$, with Y a categorical response and $X = (X_1, \ldots, X_p)$ a vector of p categorical predictors [3]. In our case the values of Y are either 0 or 1, meaning a pixel belongs to the background or it is part of a crack. The proposed method, which is called *Bayesian conditional tensor factorization* or *BCTF*, is based on the principle that a conditional probability can be expressed as a multidimensional array, i.e. a $d_1 \times \cdots \times d_p$ dimensional tensor $A = [a_{i_1,\cdots,i_p}]_{d_1\times\cdots\times d_p}$, with d_j, the number of quantization bins of the jth predictor X_j and

$$a_{i_1,\cdots,i_p} = P(Y = 1|X_1 = i_1, \cdots, X_p = i_p).$$

This *conditional probability tensor* has non-negative cells, i.e. $a_{i_1,\cdots,i_p} \geq 0$ for all $i_1, \ldots, i_p$. To determine the tensor values we perform a low rank tensor factorization, similar to *higher-order singular value decomposition* (HOSVD), introduced in [10]. Practically, computing an HOSVD of an nth-order tensor leads to the computation of n different matrix SVDs. The HOSVD decomposes the tensor $A = [a_{i_1,\cdots,i_p}]_{d_1\times\cdots\times d_p}$ as:

$$a_{i_1,\cdots,i_p} = \sum_{j_1=1}^{d_1} \cdots \sum_{j_p=1}^{d_p} s_{j_1,\cdots,j_p} \prod_{m=1}^{p} u^{(m)}_{i_m,j_m},$$

where all $U^{(m)} = [u^{(m)}_{i,j}]_{d_m\times d_m}$ are orthogonal matrices and $S = [s_{j_1,\cdots,j_p}]_{d_1\times\cdots\times d_p}$ is a so called all-orthogonal *core* tensor. In [3] it was shown that any conditional probability tensor can be decomposed in a similar (non-unique) fashion

$$\begin{aligned} &P(Y = y|X_1 = x_1, \cdots, X_p = x_p) \\ &\quad = \sum_{j_1=1}^{k_1} \cdots \sum_{j_p=1}^{k_p} \lambda_{j_1,j_2,\cdots,j_p}(y) \prod_{m=1}^{p} \pi^{(m)}_{j_m}(x_m), \end{aligned} \tag{1}$$

where moreover

$$\sum_{j_m=1}^{k_m} \pi^{(m)}_{j_m}(x_m) = 1, \tag{2}$$

holds for every combination of (m, x_m) and where we assume that the values of $k_m \in \{1, \cdots, d_m\}$ are chosen as small as possible. The factorization coefficients $\lambda_{j_1,j_2,\cdots,j_p}(y)$ can be seen as the latent class allocation probabilities and $\pi^{(m)}_{j_m}(x_m)$ as the response class probabilities and are non-negative.

A primary goal is reducing the enormous amount of covariates, as we expect that only a few features have a significant impact on the classification results, by imposing sparsity. The k_m value impacts the number of parameters used to characterize the mth predictor as well as sparsity. In the special case in which $k_m = 1$ and when taking into consideration the constraints in (2), this results in $\pi^{(m)}_1(x_m) = 1$, which means that $P(y|x_1, \cdots, x_p)$ will not depend on x_m and the mth predictor is excluded from the model. In other words, sparsity can be imposed by setting $k_m = 1$ for most m's. Additionally, k_m can be seen as the number of latent classes for the mth covariate. The levels of X_m are clustered according to their relationship with the response variable in a soft-probabilistic manner, where k_m controls the complexity of the latent structure.

To complete our Bayesian model we choose independent *Dirichlet* priors (commonly used in Bayesian statistics) for the parameters $\Lambda = \{\lambda_{j_1,\cdots,j_p};\ j_m = 1, \cdots, k_m;\ m = 1, \cdots, p\}$ and $\pi = \{\pi^{(m)}_{j_m}(x_m);\ j_m = 1, \cdots, k_m;\ x_m = 1, \cdots, d_m;\ m = 1, \cdots, p\}$,

$$\begin{aligned} &\{\lambda_{j_1,\cdots,j_p}(0), \lambda_{j_1,\cdots,j_p}(1)\} \sim \textit{Dirichlet}(1/2, 1/2), \\ &\{\pi^{(m)}_1(x_m), \cdots, \pi^{(m)}_{k_m}(x_m)\} \sim \textit{Dirichlet}(1/k_m, \cdots, 1/k_m). \end{aligned}$$

These priors have the advantages of imposing the non-negativity and summing to one constraints. The hyper parameters in the Dirichlet priors are chosen to favor placing most of the probability on a few elements, including near sparsity in these vectors. To embody our prior belief that only a small number of k_j's are greater than one we set $P(k_j = 1) = 1 - \frac{r}{p}$ and $P(k_j = k) = \frac{r}{(d_j-1)p}$ for $k = 2, \cdots, d_j,\ j = 1, \cdots, p$, and where r is the expected number of predictors to be included. To further impose sparsity, we include a restriction that $\#\{j : k_j > 1\} \leq \bar{r}$, where $\bar{r}$ is a prespecified maximum number of predictors. The effective prior on the k_j's is

$$\begin{aligned} &P(k_1 = l_1, \cdots, k_p = l_p) \\ &\quad = P(k_1 = l_1) \cdots P(k_p = l_p) I_{\{\#\{j:k_j>1\}\leq\bar{r}\}}(l_1, \cdots, l_p), \end{aligned}$$

where $I_A(\cdot)$ is the indicator function for set A. The full conditional posterior distributions of Λ and π are of the same family as the prior distributions.

5 Results & Conclusions

Fig. 1 depicts the detection results of the BCTF classifier on a detail of the Ghent Altarpiece, which is particularly challenging because it contains other objects that could be qualified as cracks.

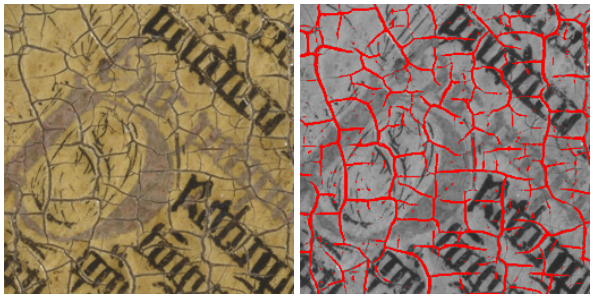

(a) Original (b) Detected cracks

Figure 1: Detection results with BCTF classifier

Next to the precise detection, the BCTF classifier has a number of appealing properties. For one, there is no parameter tweaking involved which is usually a time consuming process when working with more common image processing operations. The second interesting fact is that the output of the BCTF is actually the probability for each pixel of being a crack pixel. This means that we can propose a *probability crack map* instead of a *binary crack map* as a solution. A binary crack map is obtained by thresholding the corresponding probability map. A pixel is considered to be part of a crack when it has a crack probability of 0.5 or higher. Moreover, the BCTF method is significantly faster than the RF classifier, a patch of size 256×256 is processed within 10 seconds with BCTF while it takes 18 seconds on average with a RF classifier.

References

[1] T. Ruzic, B. Cornelis, L. Platisa, A. Pizurica, A. Dooms, W. Philips, M. Martens, M. De Mey, and I. Daubechies, "Virtual restoration of the Ghent Altarpiece using crack detection and inpainting," in *Advanced Concepts for Intelligent Vision Systems (ACIVS 2011)*, Ghent, Belgium, 2011.

[2] B. Cornelis, T. Ruzic, E. Gezels, A. Dooms, A. Pizurica, L. Platisa, J. Cornelis, M. Martens, M. D. Mey, and I. Daubechies, "Crack detection and inpainting for virtual restoration of paintings: The case of the Ghent Altarpiece," *Signal Processing*, vol. 93, no. 3, pp. 605–619, 2013.

[3] Y. Yang and D. B. Dunson, "Bayesian Conditional Tensor Factorizations for High-Dimensional Classification," Jan. 2013. [Online]. Available: http://arxiv.org/abs/1301.4950

[4] B. Cornelis, Y. Yang, J. Vogelstein, I. Daubechies, and D. Dunson, "Bayesian crack detection in ultra high resolution multimodal images of paintings," in *18th International Conference on Digital Signal Processing (DSP 2013)*, Santorini, Greece, 2013.

[5] L. Breiman, J. Friedman, R. Olshen, and C. Stone, *Classification and Regression Trees*. Monterey, CA: Wadsworth and Brooks, 1984.

[6] I. Arganda-Carreras, C. O. S. Sorzano, R. Marabini, J. M. Carazo, C. O. de Solorzano, and J. Kybic, "Consistent and elastic registration of histological sections using vector-spline regularization," in *Computer Vision Approaches to Medical Image Analysis*, ser. Lecture Notes in Computer Science, vol. 4241. Springer Berlin / Heidelberg, May 2006, pp. 85–95.

[7] A. F. Frangi, W. J. Niessen, R. M. Hoogeveen, T. van Walsum, and M. A. Viergever, "Model-based quantitation of 3-D magnetic resonance angiographic images," *IEEE Transactions on Medical Imaging*, vol. 18, no. 10, pp. 946–956, Oct. 1999.

[8] J. Serra, *Image Analysis and Mathematical Morphology*. Orlando, FL, USA: Academic Press, Inc., 1983.

[9] T. Ojala, M. Pietikäinen, and T. Mäenpää, "Multiresolution gray-scale and rotation invariant texture classification with local binary patterns," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 24, no. 7, pp. 971–987, 2002.

[10] L. D. Lathauwer, B. D. Moor, and J. Vandewalle, "A multilinear singular value decomposition," *SIAM J. Matrix Anal. Appl*, vol. 21, pp. 1253–1278, 2000.

Compressed sensing of perfusion MRI

Marie Daňková[1] and Pavel Rajmic[1]

[1]Department of Telecommunications, FEEC, Brno University of Technology

***Abstract*— Perfusion MRI is a method in medicine used mainly for diagnosing carcinoma and cardiovascular diseases. In the method a contrast agent is injected in the patient and then its concentration is observed via MRI during time. The signal(s) captured from the affected area can be described by the curve of lognormal distribution. The standard way of obtaining the measurements is very slow and does not comply with today's challenging requirements. We propose using compressed sensing to acquire much less coefficients, having minimal effect on the signal reconstruction (using the assumption that the data is a sum of low-rank matrix and sparse matrix in row spectrum).**

1 Introduction

Perfusion MRI

Perfusion imaging of the organism is an important tool in the diagnosis of cardiovascular and oncological diseases. Currently perfusion imaging may be used in a variety of imaging modalities such as magnetic resonance imaging, computed tomography, or ultrasound.

General procedure of perfusion imaging is similar for all modalities. A suitable contrast agent is incorporated into the organism as an injection or infusion and acquisition of data is performed in area of interest. Contrast agent is distributed in the body due to the cardiovascular system and its temporal and spatial distribution can be monitored.

Areas of interest may be different organs, parts of organs or individual voxels. The time course of concentration contrast agent in the region of interest is called the perfusion curve. By the analysis of perfusion curves we obtain estimates of perfusion parameters necessary for making the diagnosis.

Perfusion curve can be mathematically modeled using the density function of lognormal distribution with appropriate parameters

$$f(t) = \begin{cases} c & t \leq t_0 \\ c + \frac{S}{(t-t_0)\sigma\sqrt{2\pi}} \mathrm{e}^{-\frac{(\ln(t-t_0)-\mu)^2}{2\sigma^2}} & t > t_0 \end{cases} .$$

Here c represents the offset, t_0 is delay of bolus between injection point and area of interest, S is area under curve and μ and σ is parameters of lognormal distribution. Pefusion curve is sparse in the Fourier spectrum (Fig. 1) which we use later.

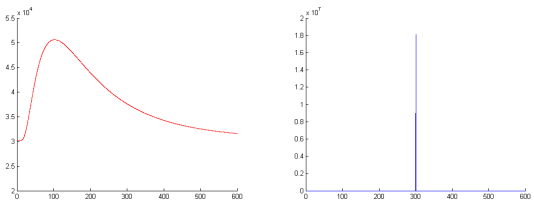

Figure 1: Example of lognormal model od perfusion MRI (left) and its amplitude spectrum (right)

Compressed sensing

Measurements in compressed sensing are linear combinations of data, and the reconstruction is nonlinear. In MRI, the signal is captured in the frequency domain, therefore a single "measurement" in this context is obtaining one Fourier coefficient. Compressed sensing helps accelerating the measurement process but this is for the price of the time for numerical reconstruction. It starts from the apriori knowledge of the signal.

2 Model specification

Preparation of data

To be able to use compressed sensing, the MRI video is reshaped into a matrix, where each column of this matrix represents a single video frame (Fig. 2). The perfusion curves form the rows of this matrix.

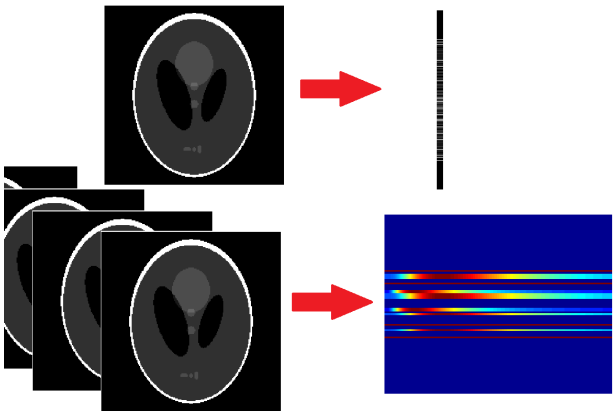

Figure 2: Reshaping video data into 2D matrix

Measurement process

Data is measured within each frame's spectrum (k-space) by application of mask (matrix consisting of ones/zeros). We use several types of mask (see Fig. 3) to measure:

- completely random with uniform distribution
- random with nonuniform distrubution
- radial (randomized or not)
- spiral (randomized or not)

Signal recovery by $L + S$ model

We use so-called L+S model for reconstruction perfusion data. In this model, the video data M is decomposed as a superposition of a low-rank matrix L (few non-zero singular values) and a sparse matrix S (few non-zero entries). See the optimization problem in Fig. 4. For better reconstruction, the Fourier coefficients which are unmeasured but have a complex-conjugated measured counterpart, are completed in this simple manner.

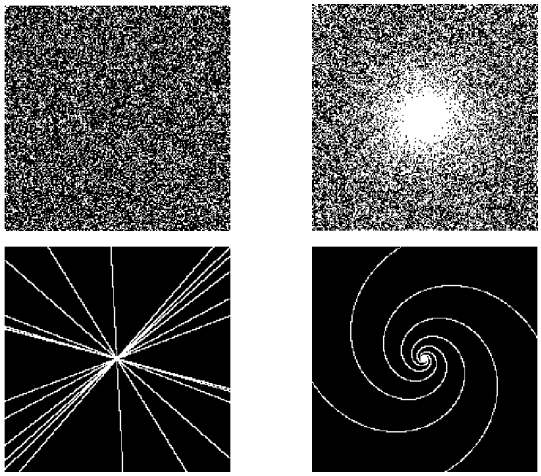

Figure 3: Measurement matrices (from top left to down right): random, nonunifom random, radial and spiral

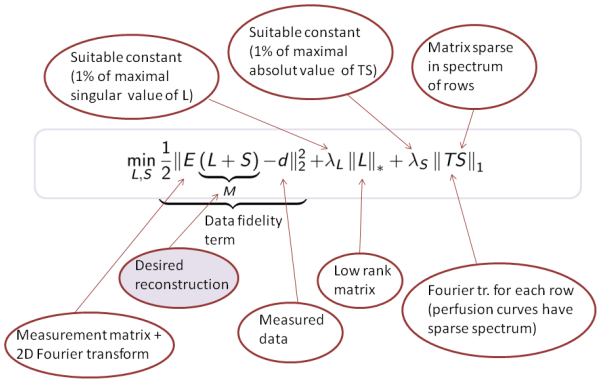
Suitable constant (1% of maximal singular value of L)
Suitable constant (1% of maximal absolut value of TS)
Matrix sparse in spectrum of rows
$\min_{L,S} \frac{1}{2}\|\underbrace{E(L+S)}_{M} - d\|_2^2 + \lambda_L \|L\|_* + \lambda_S \|TS\|_1$
Data fidelity term
Desired reconstruction
Low rank matrix
Fourier tr. for each row (perfusion curves have sparse spectrum)
Measurement matrix + 2D Fourier transform
Measured data

Figure 4: Optimization leading to the estimates L, S approximating the MRI video

Numerical solution by proximal gradient method
Proximal gradient method is for solving this type of convex minimization problem:

$$\min_x g(x) + h(x),$$

where g, h are convex and g is smooth. This problem can be solved iteratively by two steps:

$$x_k = \text{prox}_h\left(x_{k-1} - t_k \nabla g(x_{k-1})\right)$$

$$\text{prox}_h(y) = \text{argmin}_x \frac{1}{2}\|y - x\|_2^2 + h(x),$$

where prox is the so-called proximity operator. Applied to our problem, function g represents the quadratic term of the objective function and h represents the nuclear and ℓ_1 norms. Then the respective proximity operators are the soft thresholding of singular values for nuclear norm and soft thresholding of coefficients for ℓ_1 norm.

3 Results

We create phantom of perfusion data of size $100 \times 100 \times 100$ and $150 \times 150 \times 100$ pixels in Matlab. Then we measured this video using a single-type mask, but with different random parameters in each frame and we reconstructed the video by the $L + S$ algorithm. See Fig. 5. An example of the original 2D matrix, its reconstruction M and its decomposition into L and S for the random mask is depicted in Fig. 6.

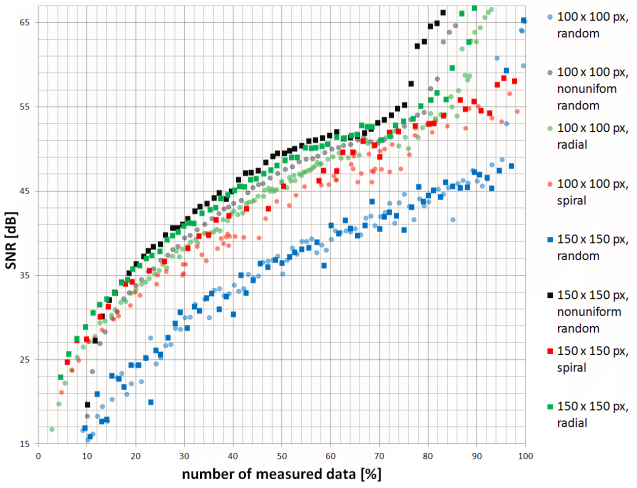
SNR [dB]
number of measured data [%]
100 x 100 px, random
100 x 100 px, nonunifom random
100 x 100 px, radial
100 x 100 px, spiral
150 x 150 px, random
150 x 150 px, nonuniform random
150 x 150 px, spiral
150 x 150 px, radial

Figure 5: Dependency of the SNR of reconstructed video on the percentage of measured data, for several types of masks and two data sizes

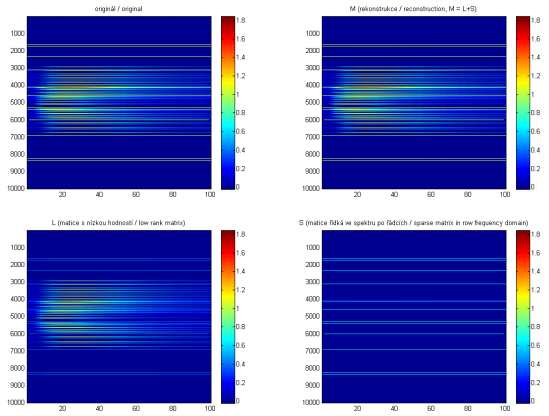

Figure 6: Original and reconstructed video data

4 Conclusion

Decomposition of the perfusion MRI matrix using the $L + S$ model can rapidly reduce number of measurements. The best performance is obtained by using the nonuniform random mask. When the data size is bigger the results are better.

References

[1] R. Otazo, E. J. Candès and D. Sodickson. Low-rank and sparse matrix decomposition for accelerated dynamic MRI with separation of background and dynamic components. Submitted to *Magnetic Resonance in Medicine*, Sept. 2013.

[2] V. Harabis, R. Kolar, M. Mezl, R. Jirik. Comparison and Evaluation of Indicator Dilution Models for Bolus of Ultrasound Contrast Agents. Physiol Meas. 2013, 34(2):151-62. doi: 10.1088/0967-3334/34/2/151

[3] P. L. Combettes, J.-C. Pesquet. Proximal Splitting Methods in Signal Processing. *Fixed-Point Algorithms for Inverse Problems in Science and Engineering*. Springer, 2011.

Compressive Hyperspectral Imaging by Out-of-Focus Modulations and Fabry-Pérot Spectral Filters

K. Degraux[1], V. Cambareri[2], B. Geelen[3], L. Jacques[1], G. Lafruit[3] and G. Setti[2].

[1]ISPGroup/ICTEAM, Université catholique de Louvain, Belgium. [2]ARCES, University of Bologna, Italy. [3]IMEC, Belgium.

***Abstract*— We describe a compressive hyperspectral imaging scheme that randomly convolves each spectral band of the data cube. This independent sensing of each wavelength relies on a tiling of Fabry-Pérot filters incorporated in the CMOS pixel grid. The compressive observations are induced by an out-of-focus spatial light modulation joined to focusing optics. While our design extends a recent monochromatic imaging scheme to the hyperspectral domain, we show that our model reaches good reconstruction performances when compared to more ideal sensing methods.**

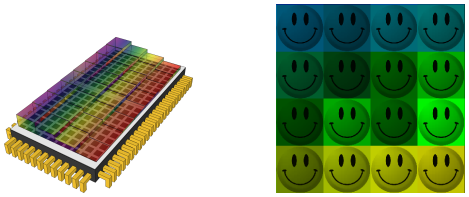

Fig. 1: (left) Tiled Fabry-Pérot filters on top of a CMOS sensor. (right) Part of a tiled representation of our low-rank and sparse HS toy example.

1 Introduction

Hyperspectral imaging is an advanced imaging technique which integrates spectroscopy into the image capturing process, providing for each pixel its light intensity as a function of wavelength. The resulting information is organized into a *hyperspectral* (HS) cube with two spatial and one spectral dimensions, which may be interpreted as a stack of images, one for each wavelength. Spectral cameras are currently used in research applications such as remote sensing [1] and food inspection [2], where this detailed spectral information may be used at application-level to identify objects and materials. Due to rapidly falling system complexity, size and cost, such technologies are expected to be adopted in more general purpose applications [3].

In order to acquire these vast amounts of data, spectral cameras typically use a time-consuming scanning approach, based on, *e.g.*, line scanning with dispersive optics (such as prisms and gratings [4]) or spectral scanning using tunable filters (*e.g.*, AOTF [5]). This problem may be overcome using snapshot acquisition, where the entire 3-D datacube is acquired during one frame period by optically *multiplexing* the contents of the 3-D cube onto a 2-D sensor. Whereas in scanning spectral imaging the data in a spectral cube is conveyed to multiple, consecutive frames, in snapshot spectral imaging this data is multiplexed onto a single frame of the sensor. However, due to the limited number of sensor pixels, snapshot imagers require a trade-off between spatial and spectral resolution.

This trade-off has been previously tackled by using *compressive sensing* (CS) techniques [6] to acquire the 3-D HS cube at sub-Nyquist sampling rates [7, 8]. In this work, we introduce a new compressive HS scheme that combines a snapshot imaging sensor based on monolithically integrated Fabry-Pérot [9] filter tiles with an optical front-end based on an out-of-focus spatial light modulator (SLM). By allowing the number of SLM patterns to vary at runtime, the system can reach higher SNR, thus operating as a hybrid setup between a scanning and a snapshot spectral camera. Since compressive sensing techniques avoid scanning, associated problems with motion blur and reduced SNR in high speed applications are also avoided.

2 Optical and Sensing Model

In this work, we extend the monochromatic compressive optical scheme developed by Björklund and Magli [10] for acquiring hyperspectral images. Our imager model is actually built around a snapshot spectral sensor, which monolithically integrates a set of Fabry-Pérot (FP) interferometers organized as tiled filters on top of a standard, off-the-shelf CMOS sensor [3] (Fig. 1). The FP filter is typically made of a transparent layer (or *cavity*) with two mirrors at each side of that layer. The cavity length and the mirror reflectivity determine the selected wavelength and the spectral bandwidth (or *full width at half maximum*) of the filter, respectively. Thanks to the monolithic integration of the filter on the sensor, the stray light in the system is heavily reduced and the sensitivity and the speed of the sensor are increased [11]. The use of CMOS process technology heavily reduces the cost and improves the compactness of the hyperspectral camera.

Let us now explain how the FP filters are combined with the compressive setup. The optical model of the imager, restricted to a 2-D section for simplicity, is illustrated in Fig. 2 (a). First, in a *monochromatic* setting, an object is imaged by conventional *objective* optics O on an *image plane*. Classical optics tells us that every point on this plane radiates light in many directions, as initially produced by the original object point. Therefore, after the image plane, the object image is replicated in many *beams* of parallel light rays. By placing an out-of-focus *spatial light modulator* (SLM) after that plane, each of such beams can be modulated differently according to its direction, before being focused (*i.e.*, "summed") on a single pixel of the detector by a second lens F. If this lens has focal length f, this pixel is located at distance $f \tan\theta$ of the optical axis. By construction, the detector actually records specific samples of the *convolution* (up to a kernel reversal) between the image and the modulation. In the case of *polychromatic* imaging with no chromatic aberrations in our optics[2], by tiling the detector with FP spectral filters we independently convolves different slices of the HS volume with the SLM modulation. As explained hereafter, by inserting a random SLM pattern, the optical scheme can compressively image the HS volume, wavelength-wise, with a partial Toeplitz sensing [12].

Mathematically, let a HS volume $\boldsymbol{x} \in \mathbb{R}^{N\times N\times L^2}$ with $N_{\text{tot}} = L^2N^2$ voxels to be acquired by the imager. The number of wavelengths is set to a square value L^2 since the detector is made of a 2-D grid of $N_{\text{s}} \times N_{\text{s}}$ pixels covered by a grid of $L \times L$ square tiles $\mathcal{T}_j$ ($1 \leqslant j \leqslant L^2$) of distinct FP filters, each selecting one wavelength λ_j. Assuming N_{s} divisible by L, each tile spectrally filters the light received by a square patch of

We thank the supercomputing facilities (CISM/UCL) and the CECI, funded by FRS-FNRS, for their computational resources. LJ and KD are funded by the FRS-FNRS.

[2]These will be fully considered in a future study.

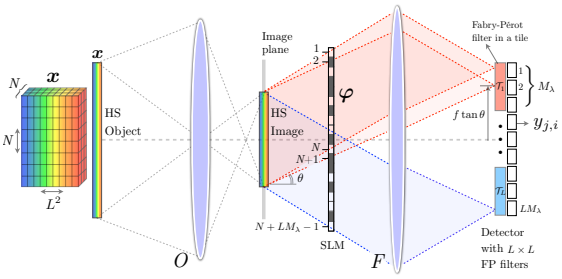

Fig. 2 (a): The compressive hyperspectral imager (1-D slice)

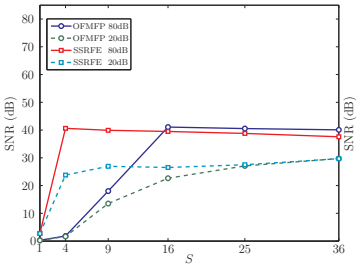

Fig. 2 (b): Low-rank/joint sparsity model.

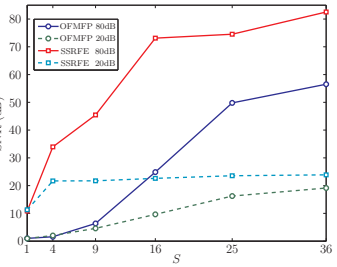

Fig. 2 (c): Total variation model.

$M_\lambda \times M_\lambda$ pixels, with $M_\lambda = N_\mathrm{s}/L$. By design (see Fig. 2 (a)), since any replication of the image must be spatially modulated at resolution N (so as to determine the spatial resolution of $\boldsymbol{x}$) and by adjusting a shift of one pixel in the detector to a shift of one pixel for the associated image beam in the SLM plane, we have $N_\mathrm{slm} = N + N_\mathrm{s} - 1 = N + LM_\lambda - 1$. For limiting the SLM resolution, we arbitrary set $N_\mathrm{slm} = 2N - 1$ so that $N_\mathrm{s} = N$ and $M_\lambda = N/L$.

Every pixel of a given tile $\mathcal{T}_j$ is influenced by a limited part $\boldsymbol{\varphi}_j \in \mathbb{R}^{(N+M_\lambda)\times(N+M_\lambda)}$ of the whole SLM modulation $\boldsymbol{\varphi} \in \mathbb{R}^{N_\mathrm{slm}\times N_\mathrm{slm}}$. Therefore, denoting by $\overline{\boldsymbol{w}}$ the *vectorization* of any matrix $\boldsymbol{w}$, the recorded light intensities $\boldsymbol{y}_j \in \mathbb{R}^{M_\lambda \times M_\lambda}$ in $\mathcal{T}_j$ read

$$\overline{\boldsymbol{y}}_j = \boldsymbol{\Phi}_j \overline{\boldsymbol{x}}_j + \overline{\boldsymbol{n}}_j \in \mathbb{R}^{M_\lambda^2}, \quad \boldsymbol{\Phi}_j := \boldsymbol{S}_j \boldsymbol{C}_j,$$

where $\boldsymbol{x}_j$ is the j^{th} slice of $\boldsymbol{x}$ at wavelength λ_j, $\boldsymbol{C}_j \in \mathbb{R}^{N^2\times N^2}$ is the Toeplitz matrix convolving an input image with $\boldsymbol{\varphi}_j$, $\boldsymbol{S}_j \in \{0,1\}^{M_\lambda^2 \times N^2}$ is a selection matrix extracting indices associated to pixels in $\mathcal{T}_j$, and $\overline{\boldsymbol{n}}_j$ accounts for possible measurement noise. Gathering all the wavelengths and vectorizing the result, the whole sensing model is $\overline{\boldsymbol{y}} = \boldsymbol{\Phi}\overline{\boldsymbol{x}} + \overline{\boldsymbol{n}}$, with $\boldsymbol{\Phi} := \mathrm{diag}(\boldsymbol{\Phi}_1, \cdots, \boldsymbol{\Phi}_{L^2})$ and $\boldsymbol{y}, \boldsymbol{n} \in \mathbb{R}^{M_\lambda \times M_\lambda \times L^2}$.

The final sensing is thus realized independently for each wavelength, as observed from the block structure of $\boldsymbol{\Phi}$. This restricts the performances of our system compared to an ideal compressive imaging device, *i.e.*, when $\boldsymbol{\Phi}$ is dense and unstructured. However, conversely to the sensing operator, the reconstruction methods recovering $\boldsymbol{x}$ from $\boldsymbol{y}$ as described in Sec. 3 exploit the 3-D structure of the data.

Compressed Sensing theory [6] shows that if a $M \times D$ sensing matrix $\boldsymbol{A}$ respects the *restricted isometry property* (RIP) with sufficiently small RIP constant, any K-*sparse* (or compressible) signals of $\mathbb{R}^D$ can be recovered (or estimated) from the possibly noisy measurement vector $\boldsymbol{Ax} \in \mathbb{R}^M$. Moreover, a structured sensing issued from a M-subsampled Toeplitz or circulant matrix whose generating entries are drawn from an iid Bernoulli ± 1 distribution, satisfies the restricted isometry property (RIP) on K-*sparse* signals if $M \geqslant M_0 = O(K^{3/2}(\log D)^{3/2})$ [12]. Since the matrices $\boldsymbol{\Phi}_j$ respect such a Toeplitz structure (even in 2-D), we can therefore expect to efficiently reconstruct each slice of the HS volume $\boldsymbol{x}$ under a sparse signal prior provided the number of observations is large. However, this work tests the regularizing capabilities of two other priors: a 3-D *total variation* (TV), for a "cartoon shape" HS model, and a low-rank and joint sparse prior [13].

We also consider a situation where the number $M_\lambda^2 = N^2/L$ of observations per band for one modulation $\boldsymbol{\varphi}$ is fixed by the detector design. An increasing number of observations is reached by merging the measurements obtained for S different random SLM patterns. The HS volume $\boldsymbol{x}$ of N^2L^2 voxels is thus observed with a total of $M_\lambda^2 L^2 S = N^2 S$ measurements, *i.e.*, we aim at reaching $1 \leqslant S \ll L^2$ with high reconstruction quality.

3 Simulations

The capability of our HS imaging scheme has been simulated as follows. We focus our study on two toy examples of size $N = 256$ and $L = 8$, *i.e.*, two structured HS volumes with $N_\mathrm{tot} = 2^{22}$ voxels. The first is a "Mondrian-like" example [15] with small 3-D *total variation* (TV) [14] made of 8 randomly generated overlapping cubes of different sizes and values in the HS volume. The second example follows a low-rank source model (LRSM) of rank $r = 4$ with (joint) sparsity level $K = N^2/16 = 4096$ [13] (see Fig. 1).

Two kinds of reconstruction methods have been used to recover these HS volumes from their compressive observations. For the "Mondrian" example, we solve $\boldsymbol{x}^* = \mathrm{argmin}_{\boldsymbol{u}} \|\boldsymbol{u}\|_{TV} + \rho\|\boldsymbol{y} - \boldsymbol{\Phi u}\|_2^2$ thanks to the TwIST algorithm [16] and by manually adjusting ρ. The recovery of the LRSM example is achieved by promoting a low-rank and joint sparse data model [13], *i.e.*,

$$\boldsymbol{x}^* = \underset{\boldsymbol{u}}{\mathrm{argmin}} \left(\sum_j \|\boldsymbol{\Psi}^T \overline{\boldsymbol{u}}_j\|_1^2\right)^{1/2} + \rho \left\|(\overline{\boldsymbol{u}}_1, \cdots, \overline{\boldsymbol{u}}_{L^2})\right\|_* \quad \text{s.t. } \|\boldsymbol{y} - \boldsymbol{\Phi u}\|_2 \leqslant \epsilon,$$

where $\|\boldsymbol{A}\|_*$ is the trace norm of $\boldsymbol{A}$, $\boldsymbol{\Psi}^T$ is the spatial Haar wavelet transform [17] at every wavelength, and the first term in the minimized cost is the ℓ_1/ℓ_2 mixed norm enforcing *joint* sparsity at every wavelength [18]. This optimization problem is solved by proximal algorithms [19, 20] with $\rho = \sqrt{K/2r}$ [13].

Fig. 2 (b) and Fig. 2 (c) display the reconstruction SNR for both methods as a function of the number of random SLM patterns S. Two noisy sensing scenarios were considered with measurement noise at 20 and 80 dB. The sensing associated to our out-of-focus modulation with FP filters (OFMFP) is confronted with the more efficient spread-spectrum random Fourier ensemble (SSRFE) sensing [21], also applied bandwise. For the two noise levels and the two considered HS examples, we see that our sensing quickly reaches similar reconstruction qualities as SSRFE for increasing S, with a stronger match for the LRSM model. Moreover, for these two toy examples and under a 20 dB noise, as soon as $S \gtrsim 16$ (*i.e.*, 25% of Nyquist rate), the SNR starts to saturate to a value related to the noise power.

4 Discussions and Conclusion

The capability of a compressive HS imager combining a CMOS detector tiled with FP filters and an out-of-focus SLM modulation has been studied. Despite an independent sensing of each wavelength, our simulations demonstrate already good performances on toy examples for a small number of modulations compared to the number of wavelengths. In a future work, the optical aberrations of the different optics and the SLM diffraction will be integrated in the sensing model. We will also study how to modulate the cube both spatially and spectrally.

References

[1] P. Shippert "Why use Hyperspectral Imagery," Photogrammetric Engineering and Remote Sensing, **70**(4): 377–380, 2004.

[2] A. A. Gowen, C.P. O'Donnell, P.J. Cullen, G. Downey and J.M. Frias, "Hyperspectral imaging - an emerging process analytical tool for food quality and safety control," Trends in Food Science and Technology, **18**(12): 590–598, 2007.

[3] N. Tack, A. Lambrechts, S. Soussan and L. Haspeslagh, "A compact, high-speed, and low-cost hyperspectral imager," Proc. SPIE 8266, Silicon Photonics VII (2012).

[4] "Imspector Fast": http://www.specim.fi/media/imspector-datasheets/imspector-fast10-ver1-11.pdf

[5] N. Gupta, R. Dahmani, K. Bennett, S. Simizu, R. Dennis and N.B. Singh, "Progress in AOTF hyperspectral imagers," Proc. SPIE Vol. 4054, 30–38 (2000).

[6] E. Candès, J. K. Romberg, T. Tao, "Stable signal recovery from incomplete and inaccurate measurements", Communications on Pure and Applied Mathematics, **59**(8):1207–1223, 2006.

[7] M.E. Gehm, R. John, D. J. Brady, R. M. Willett, and T. J. Schulz, "Single-shot compressive spectral imaging with a dual-disperser architecture," Opt. Express, **15**(21):14013–14027, 2007.

[8] H. Arguello and G. R. Arce, "Code aperture optimization for spectrally agile compressive imaging", JOSA A **28**(11):2400–2413, 2011.

[9] The Fabry-Pérot resonator, Dickmann: http://repairfaq.ece.drexel.edu/sam/MEOS/EXP03.pdf

[10] T. Björklund and Enrico Magli, "A Parallel Compressive Imaging Architecture for One-Shot Acquisition", 30th Picture Coding Symposium (PCS 2013). arXiv preprint arXiv:1311.0646 (2013).

[11] J. Loesel and D. Laubier, "Study of accessible performances of a spectro imager using a wedge filter," Proc. of SPIE, 7100, (2008).

[12] H. Rauhut, J. Romberg and J. A. Tropp, "Restricted isometries for partial random circulant matrices." Applied and Computational Harmonic Analysis **32**(2):242–254, 2012.

[13] M. Golbabaee and P. Vandergheynst, "Compressed Sensing of Simultaneous Low-Rank and Joint-Sparse Matrices", submitted to IEEE Transactions on Information Theory, p. 32, 2012.

[14] L. I. Rudin, S. Osher and E. Fatemi, "Nonlinear total variation based noise removal algorithms," Physica D: Nonlinear Phenomena **60**(1):259–268, 1992.

[15] M. J. Fadili and J-L. Starck, "Monotone operator splitting for optimization problems in sparse recovery," 16th IEEE International Conference on Image Processing (ICIP), pp.1461–1464, 2009.

[16] J. Bioucas-Dias and M. Figueiredo, "A new TwIST: two-step iterative shrinkage/thresholding algorithms for image restoration", IEEE Transactions on Image Processing, **16**(12): 2992–3004, 2007.

[17] G. Peyré, "The Numerical Tours of Signal Processing - Advanced Computational Signal and Image Processing," IEEE Computing in Science and Engineering, **13**(4):94–97, 2011.

[18] M. Kowalski, "Sparse regression using mixed norms," Applied and Computational Harmonic Analysis, **27**(3):303–324, 2009.

[19] P. L. Combettes and J.-C. Pesquet, "Proximal splitting methods in signal processing", Fixed-Point Algorithms for Inverse Problems in Science and Engineering. Springer New York, pp. 185–212, 2011.

[20] A. Chambolle and T. Pock "A first-order primal-dual algorithm for convex problems with applications to imaging" Journal of Mathematical Imaging and Vision bf 40(1):120–145, 2011.

[21] G. Puy, P. Vandergheynst, R. Gribonval, and Y. Wiaux, "Universal and efficient compressed sensing by spread spectrum and application to realistic Fourier imaging techniques," EURASIP Journal on Advances in Signal Processing, vol. 2012, pp. 1-13, 2012.

Filtered Orthogonal Matching Pursuit: Applications

Jean-François Determe† Jérôme Louveaux‡ François Horlin†

† Brussels School of Engineering (Université Libre de Bruxelles) — Avenue F. D. Roosevelt 50, Brussels, B-1050, Belgium

‡ Ecole Polytechnique de Louvain (Université Catholique de Louvain)

Email: {jdeterme, fhorlin}@ulb.ac.be, jerome.louveaux@uclouvain.be

***Abstract*— Orthogonal Matching Pursuit (OMP) algorithms enable one to recover the support of sparse vectors on the basis of indirect measurements provided that the measurement process satisfies incoherence properties. However, the problem of computing this support in a context where several measurement vectors are available is addressed to a lesser extent in the literature. This paper proposes two computationally simple approaches based on OMP to deal with these situations. Finally, the proposed algorithms are evaluated on the basis of two different problems.**

Acknowledgements

This work is funded by the Belgian *Fonds National de la Recherche Scientifique* (FNRS).

1 Introduction

This paper examines how to recover sparse vectors $\boldsymbol{\beta}[k]$ from measurements $\boldsymbol{y}_l[k]$. The inverse problem to be solved is given by Equation (1).

$$\boldsymbol{y}_l[k] = \boldsymbol{X}_l\boldsymbol{\beta}[k] + \boldsymbol{\epsilon}_l[k] \tag{1}$$

Equation (1) is often simplified when dealing with specific applied problems.

The sparse vectors $\boldsymbol{\beta}[k] \in \mathbb{R}^m$ to be estimated are assumed to remain unchanged (the dependency upon k then disappears) or to undergo limited modifications when k varies. $\boldsymbol{\epsilon}_l[k] \in \mathbb{R}^n$ is the noise added to the measurements. The entries of $\boldsymbol{\epsilon}_l[k]$ are i.i.d. random variables distributed as $\mathcal{N}(0, \sigma_\epsilon^2)$. The measurement matrix $\boldsymbol{X}_l = [\boldsymbol{x}_1\ \boldsymbol{x}_2\ \dots \boldsymbol{x}_m]$ belongs to $\mathbb{R}^{n\times m}$ with $n < m$. $\boldsymbol{x}_i \in \mathbb{R}^n$ $(1 \leq i \leq m)$ are column vectors. Index l $(1 \leq l \leq L)$ refers to different sensing matrices as measurements can be obtained from different matrices $\boldsymbol{X}_l$ within some particular frameworks.

This paper briefly describes the two proposed algorithms and then presents the envisioned applied problems used to assess their performance.

2 Exponential Decay Averaging (EDA) algorithms

Both algorithms average the output of the regular Orthogonal Matching Pursuit (OMP) algorithm [1]. *EDAOMPVote* uses the supports returned by OMP to provide an estimate of the support while *EDA-β* uses the reconstructed sparse vectors provided by OMP to produce an estimate of $\boldsymbol{\beta}$.

The developed strategies are easily extended to other reconstruction algorithms (*e.g.* ℓ_1 norm algorithms, other greedy algorithms) since the only information that is needed for the averaging procedure is the support or the reconstructed vector $\boldsymbol{\beta}$ regardless of the method that is used to obtain them.

2.1 EDAOMPVote

Initialization: `votes = vector of zeros`;
For $k = 1 \rightarrow K$ (or $l = 1 \rightarrow L$):

1. **Classical OMP algorithm:**
   Output: $\mathcal{S}$: Support of the solution
2. **Averaging procedure:**
   (a) `votes` = g∗`votes` where g is called the forgetting factor and belongs to $[0; 1] \subset \mathbb{R}$.
   (b) $\forall i \in \mathcal{S}$, `votes`($i$) = `votes`($i$) + 1
   (c) Choose the s highest entries in vector `votes`. s is assumed to satisfy $\|\boldsymbol{\beta}\|_0 \leq s$

`votes` is a vector belonging to $\mathbb{R}^m$ as there are m possible atoms to choose from.

2.2 EDA-β

Initialization: `coefficients = vector of zeros`;
For $k = 1 \rightarrow K$ (or $l = 1 \rightarrow L$):

1. **Classical OMP algorithm:**
   Output: $\boldsymbol{\beta}$ where supp($\boldsymbol{\beta}$) $= \mathcal{S}$ and $\boldsymbol{\beta}_\mathcal{S} = \text{argmin}_{\boldsymbol{\beta}}\|\boldsymbol{y} - \boldsymbol{X}_\mathcal{S}\boldsymbol{\beta}\|_2 = \boldsymbol{X}_\mathcal{S}^+\boldsymbol{y}$: Reconstructed sparse vector
2. **Averaging procedure:**
   (a) `coefficients` = g∗`coefficients` where g is the forgetting factor.
   (b) $\mathcal{S}$ = supp($\boldsymbol{\beta}$), $\forall i \in \mathcal{S}$, `coefficients`($i$) = `coefficients`($i$) + $\boldsymbol{\beta}|_i$
   (c) Choose the s highest entries in vector |`coefficients`|. s is assumed to satisfy $\|\boldsymbol{\beta}\|_0 \leq s$

---

3 Problems of interest

3.1 Noise filtering

The first problem to be addressed is that of the noise filtering described by Equation (2).

$$\boldsymbol{y}[k] = \boldsymbol{X}\boldsymbol{\beta}[k] + \boldsymbol{\epsilon}[k] \tag{2}$$

k represents a discrete time instant and $\boldsymbol{\beta}$ varies slowly over time. EDA algorithms can be successfully applied to that kind of problems although numerical simulations have revealed that simpler approaches yield better results. The forgetting factor g is usually chosen according to the rate of variation of $\boldsymbol{\beta}$.

3.2 Joint estimation

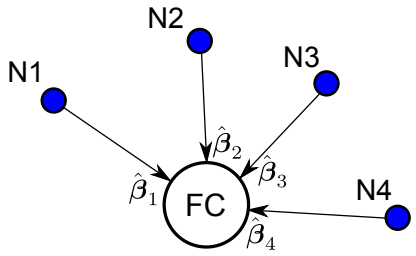

Figure 1: Joint estimation – Four nodes send data to the Fusion Center (FC)

Nodes cooperate and exchange information with the Fusion Center (FC) so as to provide an estimate of $\boldsymbol{\beta}$. The FC only uses low-complexity algorithms to perform the data fusion while the nodes perform the time-consuming operations (*i.e.* running OMP). The scalability of the system with regards to the number of nodes is therefore promoted.

The problem to be solved is given by Equation (3). The time index k has disappeared while the index l refers to the different nodes with possibly different sensing matrices $\boldsymbol{X}_l$.

$$\boldsymbol{y}_l = \boldsymbol{X}_l\boldsymbol{\beta} + \boldsymbol{\epsilon}_l \tag{3}$$

EDA algorithms (with $g = 1$) prove to be useful for solving joint estimation problems.

4 Results

4.1 Noise filtering

Numerical simulations have revealed that EDA algorithms are not well suited to Problem (2) as it is possible to obtain better results by directly filtering the measurements $\boldsymbol{y}[k]$ and then run OMP. The filter for filtering the measurements has been chosen identical to that of OMP (*i.e.* $o[k+1] = g * o[k] + u[k+1]$ where $o[k]$ and $u[k]$ are respectively the output and the input of the filter at time k).

4.2 Joint estimation

Numerical simulations have been carried out to determine how the proposed algorithms compare to centralized estimation architectures.

In centralized estimation, all the computationally intensive operations are performed by the FC and the nodes are merely sensors that transmit raw data to the FC. Hence, the performance that can be achieved is greater but the FC computational capabilities should be high and the scalability of the system is thereby limited.

The proposed simulation setup consists of several nodes with identical sensing matrices. The nodes send either estimations of the sparse vector $\boldsymbol{\beta}$ or estimations of its support depending of the algorithm that is used (*i.e.* EDA-β or EDAOMPVote).

Without entering too much into the details, the sensing matrix $\boldsymbol{X}$ belongs to $\mathbb{R}^{250\times1000}$. The sensing matrix is the same for all the simulation cases and its entries have been generated on the basis of i.i.d. Gaussian random variables.

OMP is configured to choose exactly $s = \|\boldsymbol{\beta}\|_0$ atoms and computes the estimates accordingly.

For comparison purposes, centralized OMP algorithms have been run on the same cases. Simultaneous OMP (S-OMP) and OMP Multiple Measurement Vectors (OMPMMV) are two similar centralized OMP algorithms. Detailed information about S-OMP and OMPMMV are available in [2, 3, 4]. Within the framework of centralized estimation schemes, the nodes send the raw measurements $\boldsymbol{y}_l = \boldsymbol{X}_l\boldsymbol{\beta} + \boldsymbol{\epsilon}_l$ to the FC.

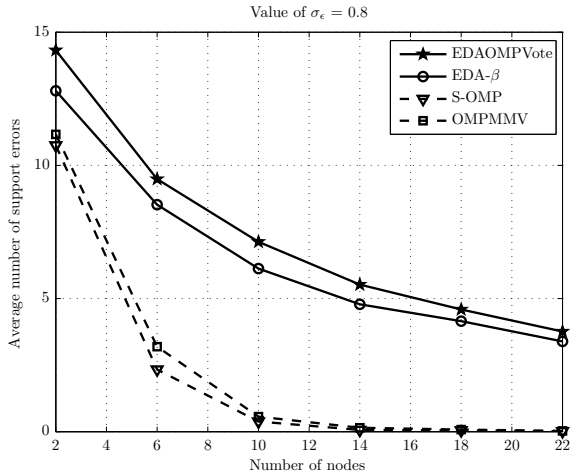

Figure 2: Joint estimation results – $\|\boldsymbol{\beta}\|_0 = 20$ – $\sigma_\epsilon = 0.8$ – $\boldsymbol{X} \in \mathbb{R}^{250\times1000}$

Figure 2 shows that S-OMP and OMPMMV deal more efficiently with high numbers of nodes although the results achieved by S-OMP/OMPMMV and EDAOMPVote/EDA-β are comparable for low numbers of nodes. Moreover, EDA-β achieves slightly better performance than EDAOMPVote.

References

[1] Cai, T. Tony, and Lie Wang. Orthogonal matching pursuit for sparse signal recovery with noise. *Information Theory, IEEE Transactions on* 57.7 (2011): 4680-4688.

[2] Chen, Jie, and Xiaoming Huo. "Theoretical results on sparse representations of multiple-measurement vectors." *Signal Processing, IEEE Transactions on* 54.12 (2006): 4634-4643.

[3] Tropp, Joel A., Anna C. Gilbert, and Martin J. Strauss. "Algorithms for simultaneous sparse approximation. Part I: Greedy pursuit." *Signal Processing* 86, no. 3 (2006): 572-588.

[4] Tropp, Joel A., Anna C. Gilbert, and Martin J. Strauss. "Simultaneous sparse approximation via greedy pursuit." In *Acoustics, Speech, and Signal Processing, 2005. Proceedings.(ICASSP'05). IEEE International Conference on*, vol. 5, pp. v-721. IEEE, 2005.

Combining sparsity and dynamics: an efficient way

Angélique Drémeau[1], Patrick Héas[2] and Cédric Herzet[2]

[1]CNRS and ESPCI ParisTech, 10 rue Vauquelin, UMR 7083 Gulliver, 75005 Paris, France.

[2]INRIA Centre Rennes - Bretagne Atlantique, Campus universitaire de Beaulieu, 35000 Rennes, France.

***Abstract*— Most linear sparse representation algorithms can be straightforwardly extended to non-linear models. Their performance however, relies on an efficient computation of the gradient of the objective function. In this paper, we focus on a particular non-linear model, defined as the nested composition of functions and propose to resort to a well-known technique from the theory of optimal control to compute the gradient. As a proof of concept, this computation is then implemented into the optimization procedure proposed by Candès *et al.*, and applied to a geophysical dynamical model.**

1 Introduction

Recent contributions have addressed the problem of exploiting sparse priors with non-linear observation models, that is

$$\mathbf{y} = h(\mathbf{x}) + \mathbf{n}, \tag{1}$$

where $h : \mathbb{R}^M \to \mathbb{R}^N$ (with $M \geq N$) is a non-linear observation operator and $\mathbf{n}$ stands for an observation noise. Extending the approach followed in the linear case, these contributions propose also a generalization of the penalty function, leading to an optimization problem of the form (or some variants thereof)

$$\hat{\mathbf{x}} = \underset{\mathbf{x}}{\operatorname{argmin}} \|\mathbf{x}\|_0 \quad \text{subject to} \quad J(\mathbf{x}) \leq \epsilon, \tag{2}$$

where $J(\mathbf{x})$ is some scalar function (*e.g.*, $J(\mathbf{x}) = \|\mathbf{y} - h(\mathbf{x})\|_2^2$) accounting for discrepancies from model (1).

Noticing that many sparse representation algorithms dealing with linear observation models rely - implicitly or explicitly - on the computation of the gradient of the function $J(\mathbf{x})$, non-linear versions of them can be straightforwardly derived. Following this idea, the extensions of the well-known algorithms IHT [1], MP [2], OMP [3] and CoSaMP [4] have thus been proposed, respectively in [5], [6], [7] and [8] .

However, whereas in the linear case, the evaluation of the gradient of $J(\mathbf{x})$ only involves multiplications by the dictionary and its transpose, its computational cost can be prohibitive in some non-linear cases. In this paper, we elaborate on this problem for the particular family of cost functions $J(\mathbf{x})$ defined as the nested composition of some functions. Formally, we write

$$J(\mathbf{x}) = \sum_{l=1}^{L} J_l \circ f_l \circ \ldots \circ f_1(\mathbf{x}), \tag{3}$$

where $\{f_l\}_{l=1}^L$ are some differentiable functions and $\circ$ stands for the function-composition operator. This type of model is for instance of interest in the ubiquitous situations where one collects partial information on the state of a *dynamical system* whose initial condition admits a sparse decomposition (see section 2.2). In particular, we emphasize that results from optimal control [9] can be exploited to provide a fast implementation of any gradient-based algorithm by taking benefit of the special structure of the non-linear model (3). We propose then a practical implementation of this computation into the optimization procedure proposed in [10].

2 Sparse Representations in Nested Non-Linear Models

In this section, we elaborate on the efficient evaluation of the gradient of $J(\mathbf{x})$ when structured as in (3). The methodology is then applied to a particular geophysical problem.

2.1 Efficient gradient computation

We use the following definitions and notations. Considering model (3), we set, $\forall l \in \{1, \ldots, L\}$, $\forall \mathbf{x} \in \mathbb{R}^M$,

$$\mathbf{s}_l \triangleq f_l \circ \ldots \circ f_1(\mathbf{x}). \tag{4}$$

We thus have $\forall l \in \{1, \ldots, L\}$,

$$\mathbf{s}_l = f_l(\mathbf{s}_{l-1}), \tag{5}$$

with the convention $\mathbf{s}_0 = \mathbf{x}$. We also define the gradient operator as

$$\nabla_{\mathbf{x}} \triangleq [\frac{\partial}{\partial x_1}, \ldots, \frac{\partial}{\partial x_M}]^T, \tag{6}$$

so that $\nabla_{\mathbf{x}}$ applied to a vector $\mathbf{z} = [z_1, ..., z_N]^T$ results in the $M \times N$ matrix whose (i, j)-th element is $\frac{\partial z_j}{\partial x_i}$.

With these notations in mind, (3) evaluated at $\mathbf{x}^\star$ can be rewritten as

$$J(\mathbf{x}^\star) = \sum_{l=1}^{L} J_l(\mathbf{s}_l^\star), \tag{7}$$

where $\mathbf{s}_l^\star$ is defined as in (4) with $\mathbf{x} = \mathbf{x}^\star$. Therefore, using the chain rule of derivative, we obtain

$$\nabla_{\mathbf{x}} J(\mathbf{x}^\star) = \sum_{l=1}^{L} \nabla_{\mathbf{x}} f_l(\mathbf{s}_{l-1}^\star)^T \, \nabla_{\mathbf{s}_l} J_l(\mathbf{s}_l^\star),$$

and from the dependence between $\mathbf{s}_l$ and $\mathbf{s}_{l-1}$,

$$\begin{aligned} \nabla_{\mathbf{x}} f_l(\mathbf{s}_{l-1}^\star)^T &= \nabla_{\mathbf{x}} f_l(f_{l-1}(\mathbf{s}_{l-2}^\star))^T, \\ &= \nabla_{\mathbf{x}} f_{l-1}(\mathbf{s}_{l-2}^\star)^T \, \nabla_{\mathbf{s}_{l-1}} f_l(\mathbf{s}_{l-1}^\star)^T. \end{aligned} \tag{8}$$

Finally, applying this expression recursively, we have

$$\nabla_{\mathbf{x}} J(\mathbf{x}^\star) = \sum_{l=1}^{L} \left(\prod_{j=1}^{l} \nabla_{\mathbf{s}_{j-1}} f_j(\mathbf{s}_{j-1}^\star)^T \right) \nabla_{\mathbf{s}_l} J_l(\mathbf{s}_l^\star). \tag{9}$$

The latter expression is exclusively based on the derivative of each function component. Its evaluation can then be performed through the following forward-backward procedure:

- The sequence $\{\mathbf{s}_l^\star\}_{l=1}^L$ is evaluated via the forward recursion

$$\mathbf{s}_l^\star = f_l(\mathbf{s}_{l-1}^\star), \quad (10)$$
$$\mathbf{s}_0^\star = \mathbf{x}^\star. \quad (11)$$

- All multiplications by a same matrix $\nabla_{\mathbf{s}_{l-1}} f_l(\mathbf{s}_{l-1}^\star)^T$ are gathered in one single operation. This is done through the backward recursion

$$\mathbf{p}_L = \nabla_{L-1} f_L(\mathbf{s}_{L-1}^\star)^T \, \nabla_{\mathbf{s}_L} J_L(\mathbf{s}_L^\star), \quad (12)$$
$$\mathbf{p}_l = \nabla_{l-1} f_l(\mathbf{s}_{l-1}^\star)^T \, (\nabla_{\mathbf{s}_l} J_l(\mathbf{s}_l^\star) + \mathbf{p}_{l+1}), \quad (13)$$

leading finally to $\mathbf{p}_0 = \nabla_{\mathbf{x}} J(\mathbf{x}^\star)$. In that way, the multiplication by each matrix $\nabla_{\mathbf{s}_{l-1}} f_l(\mathbf{s}_{l-1}^\star)^T$ is only performed once during the whole recursion.

This forward-backward procedure is widely used in geophysical applications (*e.g.*, [11]). However, to the best of our knowledge, the explicit (and motivated) use of this technique into contexts of sparsity-constrained problems has never been considered. In particular, in [12] which focuses on a similar non-linear model, this efficient computation of the gradient is not proposed.

2.2 Application to super-resolution in SQG dynamical model

As a practical example of the proposed methodology, we focus on a super-resolution problem in a geophysical context, namely the high-resolution characterization of the state of the ocean by exploiting: *i)* the Surface Quasi-Geostropic (SQG) dynamical model [13]; *ii)* a sparse prior on the initial condition of the dynamical model; *iii)* low-resolution satellite images.

We assume that the SQG model, of the form of (5), rules the evolution of some state variable $\mathbf{s}_l$. The definition of f_l depends on the considered numerical integration scheme (here a 4th-order Runge-Kunta method) but is not specified hereafter for conciseness. Moreover, as a prior knowledge, the initial state is supposed to be sparse in some redundant dictionary $\mathbf{H}$,

$$\mathbf{s}_1 = \mathbf{H}\mathbf{x}, \quad (14)$$

for some sparse vector $\mathbf{x}$. We are interested in recovering the value of $\{\mathbf{s}_l\}_{l=1}^L$ from the observation of low-dimensional images $\{\mathbf{y}_l\}_{l=1}^L$ (*e.g.*, collected by satellites), with

$$\mathbf{y}_l = \mathbf{G}_l \, \mathbf{s}_l + \mathbf{n}, \quad (15)$$

where $\mathbf{G}_l$ is some known observation matrix and $\mathbf{n}$ is an unknown corrupting noise.

In order to solve this inverse problem, we address the following optimization problem

$$\min_{\mathbf{x}} \sum_{l=1}^{L} \|\mathbf{y}_l - \mathbf{G}_l(f_l \circ \dots f_2(\mathbf{H}\mathbf{x}))\|_2^2 + \lambda \, r(\mathbf{x}), \quad (16)$$

where $\lambda > 0$ and $r(\mathbf{x}) = \sum_m \log(x_m + \epsilon)$, $\epsilon = 10^{-1}$, is some sparsity-enforcing regularizing function. A solution to (16) is searched by using the majorization-minimization optimization technique exposed in [10]; at each iteration an upper bound on the goal function is constructed by majorizing $r(\mathbf{x})$ by a weighted ℓ_1 norm. We look for the minimum of each of these majorizing functions by means of descent procedures involving the gradient of $J(\mathbf{x})$ (corresponding here to the first term in (16)) evaluated as presented in section 2.1.

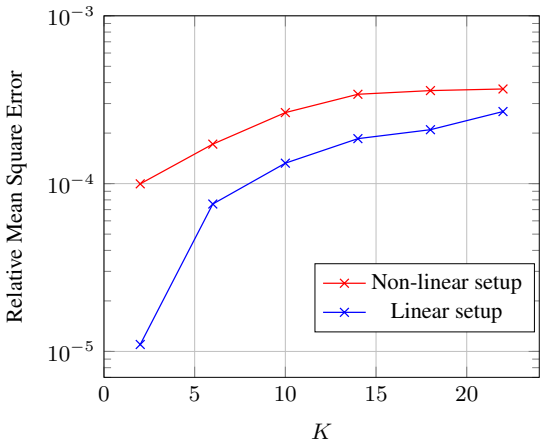

Figure 1: Relative MSE versus the number of non-zero coefficients K in the sparse vector.

Particularized to the SQG model, the evaluation of the forward-backward recursions (10)-(13) have a complexity of order $\mathcal{O}(ML)$. By comparison, using a finite-difference scheme to evaluate the gradient requires to run (at least) two forward recursions by element of $\mathbf{x}$, leading to an overall complexity of $\mathcal{O}(M^2 L)$. This order of complexity thus precludes us from using this type of approach in moderate-to-high dimensional problems.

The simulation setup considered in this paper is as follows. The state vectors $\mathbf{s}_l$ are assumed to live in 256-dimensional space. The initial condition is supposed to have a sparse decomposition in a dictionary $\mathbf{H} \in \mathbb{R}^{256\times 512}$ made up of sine and cosine functions. The observations $\mathbf{y}_l \in \mathbb{R}^{32}$ are collected at four different time instants and the observation matrices $\mathbf{G}_l$ correspond to random subsampling operators. The ratio between the number of observations and the dimension of $\mathbf{x}$ is therefore equal to $(32 \times 4)/512 = 1/4$. In Fig. 1, we represent the relative mean-square error (MSE) $||\hat{\mathbf{x}} - \mathbf{x}||_2^2 / ||\mathbf{x}||_2^2$ achieved by the minimization of (16) via the majorization-minimization procedure described above. As a point of comparison, we run the same algorithm on a linear sparse representation problem having the same problem dimensions (namely $\mathbf{y} = \mathbf{G}\mathbf{H}\mathbf{x}$ where $\mathbf{G}$ is a rate-$1/2$ random subsampling matrix). For each data point, we average the performance over 50 trials.

We can notice that the considered procedure can achieve an acceptable relative mean square error over a wide range of sparsity levels. We note also that the non-linear setup suffers from a reasonable degradation with respect to the linear setup.

3 Conclusion

In this paper, we address the problem of sparse representations in a non-linear setting. While a high computational cost of the gradient of the goal function may prevent the use of standard estimation procedures, we show that it can be overcome by applying principles from the theory of optimal control, as long as the cost function satisfies some desirable structural property. Our derivations are illustrated on a particular example dealing with the estimation of the state of a geophysical system from partial observations.

4 Acknowledgments

This work has been supported in part by the ERC under the European Union's 7th Framework Programme Grant Agreement 307087-SPARCS and by the CNRS/INSU through the LEFE funding program.

References

[1] T. Blumensath and M. E. Davies, "Iterative thresholding for sparse approximations," *Journal of Fourier Analysis and Applications*, vol. 14, no. 5-6, pp. 629–654, December 2008.

[2] S. Mallat and Z. Zhang, "Matching pursuits with time-frequency dictionaries," *IEEE Trans. On Signal Processing*, vol. 41, no. 12, pp. 3397–3415, December 1993.

[3] Y. C. Pati, R. Rezaiifar, and P. S. Krishnaprasad, "Orthogonal matching pursuit: Recursive function approximation with applications to wavelet decomposition," in *Proc. Asilomar Conference on Signals, Systems, and Computers*, 1993, pp. 40–44.

[4] D. Needell and J. A. Tropp, "Cosamp: iterative signal recovery from incomplete and inaccurate samples," *Applied and Computational Harmonic Analysis*, vol. 26, no. 3, pp. 301–321, May 2009.

[5] T. Blumensath, "Compressed sensing with nonlinear observations and related nonlinear optimization problems," *IEEE Trans. On Information Theory*, vol. 59, no. 6, pp. 3466–3474, June 2013.

[6] A. Beck and Y. C. Eldar, "Sparsity constrained nonlinear optimization: optimality conditions and algorithms," Available on arXiv:1203.4580, 2013.

[7] T. Blumensath and M. E. Davies, "Gradient pursuit for non-linear sparse signal modelling," in *European Signal Processing Conference (EUSIPCO)*, April 2008.

[8] S. Bahmani, B. Raj, and P. Boufounos, "Greedy sparsity-constrained optimization," *Journal of Machine Learning Research*, vol. 14, pp. 807–841, 2013.

[9] M. Cannon, C. Cullum, and E. Polak, *Theory of Optimal Control and Mathematical Programming*, New York, 1970.

[10] E. J. Candes, M. B. Wakin, and S. Boyd, "Enhancing sparsity by reweighted l1 minimization," *Journal of Fourier Analysis and Applications*, vol. 14, no. 5, 2008.

[11] S. O. Ba, T. Corpetti, B. Chapron, and R. Fablet, "Variational data assimilation for missing data interpolation in sst images," in *IEEE Geoscience and Remote Sensing Symposium (IGARSS)*, July 2010.

[12] A. M. Ebtehaj, M. Zupanski, G. Lerman, and E. Foufoula-Georgiou, "Variational data assimilation via sparse regularization," *EGU General Assembly*, p. 14147, 2013.

[13] J. Isern-Fontanet, G. Lapeyre, P. Klein, B. Chapron, and M. W. Hecht, "Three-dimensional reconstruction of oceanic mesoscale currents from surface information," *Journal of Geophysical Research*, vol. 113, no. C9, September 2008.

Discrete vs. Continuous Sparse Regularization

Vincent Duval[1] and Gabriel Peyré[1].

[1]CNRS and Université Paris-Dauphine.

***Abstract*— This short communication gives an overview of the main contributions from our work [6]. We put here the emphasis on the interplay between discrete (i.e. over a finite grid) and continuous (i.e. grid-free) sparse regularizations of inverse problems using ℓ^1-type regularizations. We list results ensuring that the support of the recovered sum of Diracs matches either exactly or approximately the one of the input measure when the signal-to-noise level is large enough. We also study the convergence of the discrete problem toward the continuous one when the grid size tends to zero.**

1 Introduction

We focus our attention to a class of super resolution problems, where one aims at recovering a sparse sum of Diracs from low-resolution noisy observations. This typical setup is an idealization of many problems in imaging sciences, from seismic sparse spikes deconvolution to source localization in medical imaging.

Discrete ℓ^1 regularization. We focus our attention in this article to ℓ^1-type techniques. They were initially proposed in geophysics and later studied in depth by David Donoho and coworkers [4]. A popular approach to analyze the performance of discrete ℓ^1 regularization is to assess the recovery of the positions of the non-zero coefficients, as initially introduced by Fuchs [8] and exploited by Dossal and Mallat in [5] to study sparse spike deconvolution.

Continuous total variation of measures regularization. Working over a discrete grid makes the mathematical analysis difficult. Following recent proposals (see for instance [1, 3]), we consider here sparse regularization (TV) over a continuous domain, i.e. in a grid-free setting. In this setting, the discrete ℓ^1 norm is replace by the so-called total variation of a measure (not to be confounded by the total variation of a function, which is popular to recover piecewise-constant signals and images) This allows a sharper analysis of the performances because spikes are not constrained to be a on grid, and can thus slightly shift spatially.

For deconvolution from ideal low-pass measurements, the ground-breaking paper [3] shows that this TV regularization recovers exactly sum of Diracs that are well separated. This work is further refined in [2] that studies the robustness to noise. In view of the applications of superresolution, it is crucial to understand the precise location of the recovered Diracs locations when the measurements are noisy. Partial answers to this questions are given in [7] and [1], where it is shown (under different conditions on the signal-to-noise level) that the recovered spikes are clustered tightly around the initial measure's Diracs. In this article, we fully answer the question of the position of the recovered Diracs in the setting where the signal-to-noise ratio is large enough.

Contributions We first review in Section 3 a classical result which ensures perfect support recovery over a discrete grid. We explain in Section 4 that the natural generalization of this approach in the grid free setting necessitates using a more elaborate criterion, which leads to a finer performance analysis. This is because the continuous setting allows spikes to slightly move from there initial positions, which is forbidden over a fixed grid. Lastly, Section 5 gives a precise description of the discrete-to-continuous limit, which allows a better understanding of the solution of the discrete problem when the noise is small. Note that the code to reproduce the figures of this article is available online[1]. Students and teachers might also find useful the companion numerical tour on sparse regularization over measures[2].

2 Sparse Inverse Problems Regularization

We consider the problem of recovering an input Radon measure $m_0 \in \mathcal{M}(\mathbb{T})$, where $\mathbb{T} = \mathbb{R}/\mathbb{Z}$, from noisy low resolution observations $y = \Phi(m_0) + w \in \mathrm{L}^2(\mathbb{T})$ where $w \in \mathrm{L}^2(\mathbb{T})$ is some noise and $\Phi(m)(x) = \int_{\mathbb{T}} \varphi(x, x')\mathrm{d}m(x')$ is a linear operator with a smooth kernel $\varphi \in \mathrm{C}^2(\mathbb{T} \times \mathbb{T})$. We note $\Phi^*(y)(x) = \int_{\mathbb{T}} y(x')\varphi(x', x)\mathrm{d}x'$ the adjoint defined on $\mathrm{L}^2(\mathbb{T})$. Note that for simplicity we focus our attention to 1-D inverse problems, but our results extend to arbitrary dimension.

3 Discrete Regularization

A classical way to regularize the inverse problem is to use ℓ_1 regularization over a discrete grid $z = (z_i)_{i=1}^N \subset \mathbb{T}^N$ of N points, and solve the finite dimensional inverse problem

$$\min_{a \in \mathbb{R}^N} \frac{1}{2}\|\Phi_z a - y\|^2 + \lambda\|a\|_1 \quad \text{where} \quad \|a\|_1 = \sum_{i=1}^N |a_i|, \quad (1)$$

where $\lambda > 0$ should be adapted to the noise level, where $\Phi_z a = \sum_i \varphi(\cdot, z_i)a_i$ is the restriction of the operator to measures supported on the grid z. For a solution $a^\star$ of (1), one thus obtains an estimate $m_{z,a^\star} = \sum_{i=1}^N a_i^\star \delta_{z_i}$ of the initial measure m_0.

We assume that the input measure $m_0 = m_{z,a_0}$ is supported on the grid, with a support $I = \mathrm{Supp}(a_0)$ of non-zero indexes. The following classical theorem shows that if the following precertificate (we assume Φ_I injective)

$$\eta_{\mathrm{F}} = \Phi^* \Phi_I^{*,+} \mathrm{sign}(a_{0,I}) \quad \text{where} \quad \Phi_I = (\varphi(\cdot, x_i))_{i \in I}. \quad (2)$$

does not saturate on the grid outside the support, then the discrete regularization recovers the correct support. Here $a_{0,I} = (a_{0,i})_{i \in I}$ is the sub-vector indexed by I and $\Phi_I^{*,+}$ denotes the pseudo-inverse of Φ_I^*.

[1] https://github.com/gpeyre/2013-FOCM-SparseSpikes/

[2] Accessible from http://www.numerical-tours.com

Theorem 1 (Fuchs [8]). *If* $\forall\, j \notin I, |\eta_{\mathrm{F}}(x_j)| < 1$, *then for* $\|w\|$ *small enough and* $\lambda \sim \|w\|$, *the solution* a_λ *of* (1) *is unique and* $\mathrm{Supp}(a_\lambda) = \mathrm{Supp}(a_0)$.

4 Continuous Regularization

A major issue is that Theorem 1 is overly pessimistic in practical applications, because recovering the exact positions of the spikes is not possible if the grid size is too small. One can indeed show that the condition $|\eta_{\mathrm{F}}(x_j)| < 1$ is always violated for a grid fine enough. To allow for a finer analysis, one can replace (1) by its counterpart over the space of measures

$$\min_{m\in\mathcal{M}(\mathbb{T})} \frac{1}{2}\|\Phi(m)-y\|^2+\lambda\|m\|_{\mathrm{TV}}, \quad \|m\|_{\mathrm{TV}} = \sup_{\|g\|_\infty\leqslant 1} \int g\mathrm{d}m \tag{3}$$

where the supremum is taken over continuous functions. Note that for a discrete measure, one has $\|m_{z,a}\|_{\mathrm{TV}} = \|a\|_1$.

We now assume an arbitrary input measure $m_0 = m_{x_0,a_0}$ where $x_0 \in \mathbb{T}^s$. While the discrete pre-certificate η_{F} defined in (2) is only constrained to interpolate the sign of a_0 on the support, working over a grid-free domain enables to use refined constraints. The following Theorem 2 shows that support recovery is indeed governed by a pre-certificate which interpolates the sign of the input measure together with vanishing derivatives at the support locations

$$\eta_{\mathrm{V}} = \Phi^*\Gamma_{x_0}^{+,*}\left(\mathrm{sign}(a_0),0\right)^* \in \mathrm{L}^2(\mathbb{T})$$

where we have introduced the following operator (assumed to be injective)

$$\Gamma_x : (a,b) \in \mathbb{R}^{2s} \mapsto \sum_{i=1}^{s} a_i\varphi(\cdot,x_i) + b_i\varphi'(\cdot,x_i) \in \mathrm{L}^2(\mathbb{T}).$$

Theorem 2 (Duval-Peyré [6]). *If* $\forall\, i, \eta_{\mathrm{V}}''(x_{0,i}) \neq 0$ *and* $\forall\, t \notin x_0, |\eta(t)| < 1$, *then for* $\|w\|$ *small enough and* $\lambda \sim \|w\|$, *the solution* $m_\lambda = m_{x_\lambda,a_\lambda}$ *of* (3) *is unique, is composed of* s *Diracs, and* $\|a_\lambda - a_0\|_\infty = O(\|w\|), \|x_\lambda - x_0\|_\infty = O(\|w\|)$.

The important point is that (contrary to η_{F}) since η_{V} has vanishing derivatives on the support, the non-saturation constraint (i.e. $|\eta| < 1$ outside the support) is much more likely to hold in the continuous setting (i.e. $\eta = \eta_{\mathrm{V}}$) than in the discrete one (i.e. $\eta = \eta_{\mathrm{F}}$). This is illustrated by Figure 1, on the case where $\varphi(x,x') = \sin((2f_c+1)\pi(x-x'))/\sin(\pi(x-x'))$ is an ideal low-pass filter with cut-off frequency $f_c = 6$.

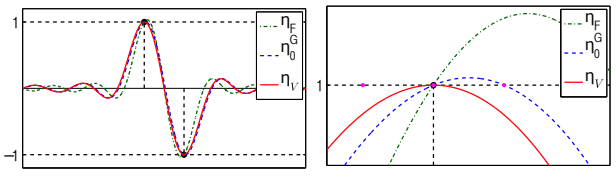

Figure 1: Comparison of the pre-certificates η_{F} and η_{V}. The figure on the right is a zoom of the left one near the left spike. The Fuchs pre-certificate η_{F} (dash-dot green line) is above 1 at some points of the grid: Theorem 1 does not apply and the support is not stable for the discrete problem. This contrasts with η_{V} which is strictly below one outside the support of the input measure: Theorem 2 thus does apply and the support is stable for the continuous problem. The certificate $\eta_0^{\mathcal{G}}$ (dashed blue line) governs the support at low noise regime on the discrete grid (see Theorem 3 and [6]).

5 From Discrete to Continuous

The following results shows that if the continuous problem has a stable support (i.e. Theorem 2 can be applied) then the solution of the discrete problem is supported on pairs of Diracs surrounding the initial measure. For simplicity, we suppose that the considered grids z are uniform.

Theorem 3 (Duval-Peyré [6]). *Let* $m_0 = m_{x_0,a_0}$ *be such that* $x_0 \subset z$ *is on the discretization grid, i.e.* $x_0 = z_I$ *for some index set* I. *We suppose that the hypotheses of Theorem 2 hold for* m_0. *Then if* N *is large enough,* $\|w\|$ *small enough and* $\lambda \sim \|w\|$, *the solution of* a_λ *of* (1) *is unique and* $\mathrm{Supp}(a_\lambda) \subset I \cup \tilde{I}$ *where* $\tilde{I} = (i+\varepsilon_i)_{i\in I}$ *with* $\varepsilon_i \in \{-1,+1\}$.

This result is illustrated on Figure 2, in the case of an ideal low pass filter with no noise. It shows that for $\lambda = 0$, the discrete problem recovers exactly the input measure, but as λ increases (and similarly if some small noise w is added), two spurious spikes, located on neighboring grid points indexed by $\tilde{I}$, pollutes the recovered measures. This duplication of spikes is linked to the inability of the discretized representation of measure to cope with translation of the Diracs, as this is achieved by the continuous one studied in Section 4. At the same time, this gives a precise geometrical understanding of the convergence of the discrete problem toward the continuous one (duplicated spikes become closer and closer to the input spikes as N increases).

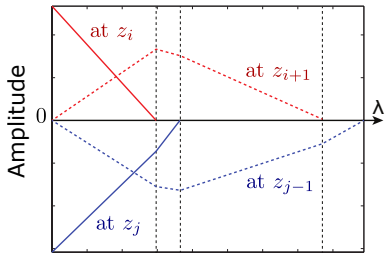

Figure 2: Display of the solution path (as a function of λ) for the measure displayed on Figure 1. This shows the amplitudes of the coefficients at z_i, resp. z_j, (continuous line) and at the next, resp. previous, point of the grid (dashed line) as λ varies.

References

[1] J-M. Azais, Y. De Castro, and F. Gamboa. Spike detection from inaccurate samplings. Technical report, January 2013.

[2] E. J. Candès and C. Fernandez-Granda. Super-resolution from noisy data. *To appear in Journal of Fourier Analysis and Applications*, 2013.

[3] E. J. Candès and C. Fernandez-Granda. Towards a mathematical theory of super-resolution. *Communications on Pure and Applied Mathematics. To appear.*, 2013.

[4] S.S. Chen, D.L. Donoho, and M.A. Saunders. Atomic decomposition by basis pursuit. *SIAM journal on scientific computing*, 20(1):33–61, 1999.

[5] C. Dossal and S. Mallat. Sparse spike deconvolution with minimum scale. In *Proceedings of SPARS*, pages 123–126, November 2005.

[6] V. Duval and G. Peyré. Exact support recovery for sparse spikes deconvolution. Technical report, Preprint hal-00839635, 2013.

[7] C. Fernandez-Granda. Support detection in super-resolution. *Proc. Proceedings of the 10th International Conference on Sampling Theory and Applications*, pages 145–148, 2013.

[8] J.J. Fuchs. On sparse representations in arbitrary redundant bases. *IEEE Transactions on Information Theory*, 50(6):1341–1344, 2004.

Dictionary learning for efficient classification based on soft-thresholding

Alhussein Fawzi[1], Mike Davies[2] and Pascal Frossard[1].

[1]Signal Processing Laboratory (LTS4), EPFL. [2]IDCoM, The University of Edinburgh.

***Abstract*— We consider in this paper an efficient classification scheme built by applying a soft-thresholding non linear map followed by a linear classifier. We propose a supervised dictionary learning algorithm tailored for this classification architecture. We cast the dictionary learning problem, that learns jointly the discriminative dictionary and the hyperplane, as a difference of convex (DC) problem and solve it efficiently with DCA. We show on texture classification experiments that our method compares favorably to standard classifiers in terms of accuracy and computational complexity.**

1 Introduction

Object classification is one of the most challenging tasks in computer vision. Recently, techniques relying on sparse coding for feature extraction have led to state of the art results on standard datasets [1, 2, 3, 4]. However, the prohibitive cost of computing the non-linear sparse representation of an image is a major drawback that limits the applicability of such techniques to large-scale vision problems. Indeed, the sparse representation in a dictionary $\mathbf{D}$ of an image $\mathbf{x}$

$$\underset{\mathbf{c}}{\operatorname{argmin}} \|\mathbf{x} - \mathbf{D}\mathbf{c}\|_2^2 + \lambda\|\mathbf{c}\|_1, \tag{1}$$

involves a non trivial optimization problem. The development of efficient techniques for solving (or approximating the solution of) Eq. (1) has been the focus of much recent research [5, 6, 7, 8]. In this work, we consider an efficient and simple alternative to the sparse coding mapping of Eq. (1). Specifically, we consider the *soft-thresholding*[1] non-linear feature extraction mapping defined by

$$(\mathbf{D}^\mathbf{T}\mathbf{x} - \alpha)_+, \tag{2}$$

where $(\cdot)_+ = \max(0, \cdot)$ and α is a parameter that controls the sparsity of the feature vectors. The usefulness of this nonlinear mapping for classification has already been observed [3] and is popular in a number of deep learning architectures due to its efficiency [9, 10]. We focus here on a classification scheme where features (obtained with the soft thresholding mapping) are fed to a linear classifier $\mathbf{w}$. We design a *supervised dictionary learning* algorithm tailored for this classification architecture. Specifically, we learn *jointly* a discriminative dictionary $\mathbf{D}$ and a hyperplane $\mathbf{w}$ that separates the two classes in the feature space. We illustrate the effectiveness of our approach in terms of classification accuracy and computational efficiency on a texture classification task.

2 Problem formulation

We focus in this paper on a binary classification problem. Let $\mathbf{X} = [\mathbf{x_1}|\dots|\mathbf{x_m}] \in \mathbb{R}^{n\times m}$ and $\mathbf{y} = [y_1|\dots|y_m] \in \{-1, 1\}^m$ denote the set of training points and their associated labels. We consider the following supervised dictionary learning formulation, that optimizes jointly for the dictionary $\mathbf{D} \in \mathbb{R}^{n\times N}$ and the linear classifier $\mathbf{w} \in \mathbb{R}^N$

$$\underset{\mathbf{D},\mathbf{w}}{\operatorname{argmin}} \sum_{i=1}^{m} L(y_i \mathbf{w}^\mathbf{T}(\mathbf{D}^\mathbf{T}\mathbf{x_i} - \alpha)_+) + \frac{\nu}{2}\|\mathbf{w}\|_2^2, \tag{3}$$

where L denotes a convex loss function that penalizes incorrect classification of a training sample and ν is a regularization parameter which prevents overfitting. The above optimization problem attempts to find a dictionary $\mathbf{D}$ and a linear separator $\mathbf{w}$ such that $\mathbf{w}^\mathbf{T}(\mathbf{D}^\mathbf{T}\mathbf{x_i} - \alpha)_+$ has the same sign as y_i on the training set. We adopt the hinge loss function, in this paper, which is defined by $L(x) = \max(0, 1 - x)$.

The problem in Eq. (3) is difficult to solve in general. We use a suitable change of variables to put it in a nice form (difference of convex functions) that provides a practical and efficient solution. Specifically, using the change of variables $\mathbf{u_j} \leftarrow |w_j|\mathbf{d_j}$, $v_j \leftarrow |w_j|$ and $s_j \leftarrow \operatorname{sgn}(w_j)$, the problem in Eq.(3) can be rewritten[2] in the following way:

$$\underset{\mathbf{U},\mathbf{v},\mathbf{s}}{\operatorname{argmin}} \sum_{i=1}^{m} L\left(y_i \sum_{j=1}^{N} s_j(\mathbf{u_j}^\mathbf{T}\mathbf{x_i} - \alpha v_j)_+\right) + \frac{\nu}{2}\|\mathbf{v}\|_2^2, \tag{4}$$
$$\text{subject to } \mathbf{v} > 0.$$

Note that s_j (i.e., the sign of w_j) essentially encodes the "class" of atom $\mathbf{d_j}$. In other words, an atom $\mathbf{d_j}$ for which $s_j = +1$ is most likely to be active in class 1 samples. Conversely, atoms that satisfy $s_j = -1$ are most likely active for class -1 samples. We assume here that the vector $\mathbf{s}$ is known a priori. In other words, this means that we have a prior knowledge on the proportion of class 1 to class -1 atoms in the dictionary. For example, setting half of the entries of the vector $\mathbf{s}$ to be equal to $+1$ and the other half to -1 encodes the prior knowledge that we are searching for a dictionary with balanced class-specific atoms.

Moreover, for the purpose of optimization, we approximate the term $(\mathbf{u_j}^\mathbf{T}\mathbf{x_i} - \alpha v_j)_+$ in Eq.(4) with a smooth function $q(\mathbf{u_j}^\mathbf{T}\mathbf{x_i} - \alpha v_j)$ where $q(x) = \frac{1}{\beta}\log(1 + \exp(\beta x))$ and β is a parameter that controls the accuracy of the approximation. Note that this approximation is done only to make the optimization easier at the dictionary learning stage. At test time, soft-thresholding is applied for feature extraction.

Thus, we end up with the optimization problem (P)[3]:

$$\text{(P)}: \quad \underset{\mathbf{U},\mathbf{v}}{\operatorname{argmin}} \quad \sum_{i=1}^{m} L\left(y_i \sum_{j=1}^{N} s_j q(\mathbf{u_j}^\mathbf{T}\mathbf{x_i} - \alpha v_j)\right) + \frac{\nu}{2}\|\mathbf{v}\|_2^2,$$
$$\text{subject to } \mathbf{v} \geq \epsilon.$$

[1]Soft-thresholding is usually defined as $\operatorname{sgn}(x)(|x| - \alpha)$. We consider only the positive part here.

[2]The two problem formulations are equivalent only when the components of $\mathbf{w}$ are restricted to be all non zero. This assumption is not limiting as zero components in the optimal hyperplane of Eq. (3) can be removed, resulting in a smaller dictionary.

[3]The strict inequality $\mathbf{v} > 0$ is replaced with $\mathbf{v} \geq \epsilon$, where ϵ is a small positive constant number. The latter constraint is easier to handle for optimization, yet both constraints are essentially equivalent in practice.

Table 1: Classification accuracy of different methods for texture classification, for different dictionary sizes.

Classification task	*pressedcl* vs. *pigskin*				*d4* vs. *d5*				*bark* vs. *woodgrain*				*paper* vs. *peb54*				*ice* vs. *grass*			
Dictionary size	10	50	100	400	10	50	100	400	10	50	100	400	10	50	100	400	10	50	100	400
Linear classifier [%]	49				50				50				53				46			
Sparse coding [%]	52	73	79	84	60	70	73	**75**	66	**91**	**94**	**94**	58	**75**	**77**	**80**	53	56	65	68
Proposed approach [%]	**72**	**78**	**80**	**86**	**68**	**72**	**74**	**75**	**80**	90	91	89	**71**	73	75	78	**71**	**76**	**78**	**75**
Nearest neighbour [%]	65	74	75	72	64	64	63	56	74	84	86	86	64	65	64	55	59	61	59	50

Algorithm 1 DCA algorithm for solving (P)

1. Choose any initial point: $\mathbf{U^0}$ and $\mathbf{v^0} \geq \epsilon$.

2. For $k \in \mathbb{N}$,

2.1 Compute $(\mathbf{A}, \mathbf{b}) = \nabla h(\mathbf{U^k}, \mathbf{v^k})$.

2.2 Solve the convex optimization problem:

$$(\mathbf{U^{k+1}}, \mathbf{v^{k+1}}) \leftarrow \underset{(\mathbf{U},\mathbf{v})}{\operatorname{argmin}} \{g(\mathbf{U}, \mathbf{v}) - Tr(\mathbf{U}^T\mathbf{A}) - \mathbf{v}^T\mathbf{b}\}$$
$$\text{subject to } \mathbf{v} \geq \epsilon.$$

2.3 If $(\mathbf{U^{k+1}}, \mathbf{v^{k+1}}) \approx (\mathbf{U^k}, \mathbf{v^k})$, return $(\mathbf{U^{k+1}}, \mathbf{v^{k+1}})$.

Given the optimal $(\mathbf{U}, \mathbf{v})$, $\mathbf{D}$ and $\mathbf{w}$ can be obtained by using the previously given change of variables. A test point $\mathbf{x}$ is then assigned to label $+1$ if $\mathbf{w^T}(\mathbf{D^T x} - \alpha)_+ > 0$, and label -1 otherwise.

3 DC optimization

Problem (P) is typically a nonconvex optimization problem that is hard to solve using traditional methods, such as gradient descent or Newton-type methods. We show however that the problem (P) can be written as a *difference of convex* (DC) program [11]. A DC problem is an optimization problem which consists in minimizing a function of the form $g - h$, where both g and h are convex, subject to difference of convex constraints. DC problems are well studied optimization problems and efficient optimization algorithms have been proposed [11, 12] and have been shown to work well in practice (see [13] and references therein).

Proposition 1. *For any convex loss function L and any convex function q, the problem (P) is DC. Moreover, when $L(x) = \max(0, 1-x)$, the objective function of problem (P) is equal to $g - h$ with:*

$$g = \frac{\nu}{2}\|\mathbf{v}\|_2^2 + \sum_{i=1}^{m} \max\Big(\sum_{j:s_j=y_i} q(\mathbf{u_j^T x_i} - \alpha v_j), 1 + \sum_{j:s_j\neq y_i} q(\mathbf{u_j^T x_i} - \alpha v_j) \Big)$$

$$h = \sum_{i=1}^{m} \sum_{j:s_j=y_i} q(\mathbf{u_j^T x_i} - \alpha v_j).$$

We solve efficiently problem (P) for a local optimal solution using the DC minimization algorithm (DCA) [12] presented in Algorithm 1. In a nutshell, each iteration of DCA solves a convex optimization problem obtained by linearizing h around its current point. Although DCA is only guaranteed to reach a local minima, the authors of [12] state that DCA often converges to a global optimum.

Table 2: Time taken [ms] to compute feature vectors for sparse coding and soft thresholding.

Dictionary size	10	50	100	400
Sparse coding [ms]	9	41	110	375
Soft thresholding [ms]	1	2	3	11

4 Experimental results

We now illustrate the effectiveness of our approach on texture image classification experiments. We consider texture images from the 32 Brodatz dataset[4]. For each pair of textures under test, we build the training set by randomly selecting 200 patches per texture of size 12×12 from the left half of the images, and the test data is constructed by selecting 1000 patches per image from the right half of the images.

For our discriminative dictionary learning approach, we fix half of the entries of $\mathbf{s}$ to be equal to $+1$ and the other half to be equal to -1. We set moreover $\alpha = 1$, $\beta = 50$ and the regularization parameter ν to 0. We initialize half of the vectors in $\mathbf{U}$ with random patches from the training set of class 1, and the other half with random patches from the training set of class 2. We also set $\mathbf{v^0}$ to $\mathbf{1}$. At test time, a datapoint $\mathbf{x}$ is assigned to class 1 if $\mathbf{w^T}(\mathbf{D^T x} - \alpha)_+ > 0$ and class -1 otherwise.

We compare our approach to several classification algorithms. The first of these methods is a simple linear SVM classifier. Also, we consider a sparse representation based classifier similar to that of [2], where a dictionary is first learned (in a reconstructive manner) that sparsifies the training data[5]. A linear SVM is then trained on the sparse codes. At test time, the classification is performed by solving Eq. (1) on the test image, and then applying a linear classifier. Finally, the last comparative scheme is a nearest neighbour classifier. The dictionary is built by applying a K-Means clustering on the training samples of each class separately. At test time, the sample is assigned the label of the dictionary atom (i.e., cluster) that is closest to it.

Table 1 compares these methods on 5 binary classification tasks, with different dictionary sizes. Our proposed approach clearly outperforms linear and nearest neigbour classifiers in all tasks. Moreover, our method compares favorably to sparse coding for small dictionary sizes. For larger dictionaries, sparse coding yields in some cases a better classification accuracy. This is due to the fact that the dictionary learning step in classification tasks is crucial when the dictionary is constrained to be small [3]. For larger dictionaries, there seems to be less performance gap between reconstructive and discriminative approaches for learning $\mathbf{D}$. As shown in Table 2, soft-thresholding based features are much more efficient to compute than sparse codes, which makes our method particularly suited to large-scale problems or when computational power is limited. In a longer version of this work (in preparation), we provide more experimental results and compare the proposed approach to other algorithms.

[4] `http://www.cse.oulu.fi/CMV/TextureClassification`

[5] We use the SPAMS toolbox for dictionary learning and sparse representation [14].

References

[1] John Wright, Allen Yang, Arvind Ganesh, Shankar Sastry, and Yi Ma, "Robust face recognition via sparse representation," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 31, no. 2, pp. 210–227, 2009.

[2] Rajat Raina, Alexis Battle, Honglak Lee, Benjamin Packer, and Andrew Y Ng, "Self-taught learning: transfer learning from unlabeled data," in *International Conference on Machine Learning (ICML)*. ACM, 2007, pp. 759–766.

[3] Adam Coates and Andrew Y Ng, "The importance of encoding versus training with sparse coding and vector quantization," in *International Conference on Machine Learning (ICML)*, 2011, vol. 8, p. 10.

[4] Julien Mairal, Francis Bach, and Jean Ponce, "Task-driven dictionary learning," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 34, no. 4, pp. 791–804, 2012.

[5] Karol Gregor and Yann LeCun, "Learning fast approximations of sparse coding," in *International Conference on Machine Learning (ICML)*, 2010.

[6] Alex Bronstein, Pablo Sprechmann, and Guillermo Sapiro, "Learning efficient structured sparse models," *International Conference on Machine Learning (ICML)*, 2012.

[7] Tong Tong Wu and Kenneth Lange, "Coordinate descent algorithms for lasso penalized regression," *The Annals of Applied Statistics*, pp. 224–244, 2008.

[8] Koray Kavukcuoglu, Marc'Aurelio Ranzato, and Yann LeCun, "Fast inference in sparse coding algorithms with applications to object recognition," *arXiv preprint arXiv:1010.3467*, 2010.

[9] Koray Kavukcuoglu, Pierre Sermanet, Y-Lan Boureau, Karol Gregor, Michaël Mathieu, and Yann LeCun, "Learning convolutional feature hierarchies for visual recognition," *Advances in Neural Information Processing Systems (NIPS)*, pp. 1090–1098, 2010.

[10] Kevin Jarrett, Koray Kavukcuoglu, MarcAurelio Ranzato, and Yann LeCun, "What is the best multi-stage architecture for object recognition?," in *International Conference on Computer Vision (ICCV)*, 2009, pp. 2146–2153.

[11] Reiner Horst, *Introduction to global optimization*, Springer, 2000.

[12] Pham Dinh Tao and Le Thi Hoai An, "A dc optimization algorithm for solving the trust-region subproblem," *SIAM Journal on Optimization*, vol. 8, no. 2, pp. 476–505, 1998.

[13] Pham Dinh Tao et al., "The dc (difference of convex functions) programming and dca revisited with dc models of real world nonconvex optimization problems," *Annals of Operations Research*, vol. 133, no. 1-4, pp. 23–46, 2005.

[14] Julien Mairal, Francis Bach, Jean Ponce, and Guillermo Sapiro, "Online learning for matrix factorization and sparse coding," *The Journal of Machine Learning Research*, vol. 11, pp. 19–60, 2010.

Semidefinite Programming Based Preconditioning for More Robust Near-Separable Nonnegative Matrix Factorization

Nicolas Gillis[1] and Stephen A. Vavasis[2].

[1]Department of Mathematics and Operational Research, Université de Mons

Rue de Houdain 9, 7000 Mons, Belgium – Email: nicolas.gillis@umons.ac.be

[2]Department of Combinatorics and Optimization, University of Waterloo

Waterloo, Ontario N2L 3G1, Canada – Email: vavasis@math.uwaterloo.ca.

***Abstract*— Nonnegative matrix factorization (NMF) under the separability assumption can provably be solved efficiently, even in the presence of noise, and has been shown to be a powerful technique in document classification and hyperspectral unmixing. This problem is referred to as near-separable NMF and requires that there exists a cone spanned by a small subset of the columns of the input nonnegative matrix approximately containing all columns. In this talk, we present a preconditioning based on semidefinite programming making the input matrix well-conditioned. This in turn can improve significantly the performance of near-separable NMF algorithms which is illustrated on the popular successive projection algorithm (SPA). The new preconditioned SPA is provably more robust to noise, and outperforms SPA on several synthetic data sets.**

1 Introduction

Nonnegative matrix factorization (NMF) is a powerful dimensionality reduction technique as it automatically extracts sparse and meaningful features from a set of nonnegative data vectors. Given n nonnegative m-dimensional vectors gathered in a nonnegative matrix $M \in \mathbb{R}^{m\times n}_+$ and a factorization rank r, NMF computes two nonnegative matrices $W \in \mathbb{R}^{m\times r}_+$ and $H \in \mathbb{R}^{r\times n}_+$ such that $M \approx WH$. Unfortunately, NMF is NP-hard in general [1]. However, if the input data matrix M is r-separable, that is, if there exists an index set $\mathcal{K}$ with cardinality r and an r-by-n nonnegative matrix H such that $M = M(:,\mathcal{K})H$, then the problem can be solved in polynomial time [2]. In the presence of noise, this problem is referred to as near-separable NMF. Although the separability condition is rather strong, it makes sense in several applications, e.g., in text mining [3], hyperspectral unmixing [4] and blind source separation [5]. Note that this problem is equivalent to finding a row-sparse matrix X such that $||M - MX||_F^2$ is minimized [6, 7].

In this talk, we present the results from [8]: we show how to make near-separable NMF algorithms more robust using preconditioning. We focus on a simple yet effective near-separable NMF algorithm, namely, the successive projection algorithm.

2 Successive Projection Algorithm

The successive projection algorithm is a simple but fast and robust recursive algorithm for solving near-separable NMF; see Algorithm SPA. At each step of the algorithm, the column of the input matrix $\tilde{M}$ with maximum ℓ_2 norm is selected, and then $\tilde{M}$ is updated by projecting each column onto the orthogonal complement of the columns selected so far. It was first introduced in [9], and later proved to be robust in [10].

Algorithm SPA Successive Projection Algorithm [9, 10]

Input: Near-separable matrix $\tilde{M}$, rank r.
Output: Set of r indices $\mathcal{K}$ such that $\tilde{M}(:,\mathcal{K}) \approx W$.

1: Let $R^{(1)} = \tilde{M}$, $\mathcal{K}^{(1)} = \{\}$.
2: **for** $k = 1 : r$ **do**
3: $\quad p = \operatorname{argmax}_j ||R^{(k)}_{:j}||_2$.
4: $\quad R^{(k+1)} = \left(I - \frac{R^{(k)}_{:p} R^{(k)\,T}_{:p}}{||R^{(k)}_{:p}||_2^2}\right) R^{(k)}$.
5: $\quad \mathcal{K}^{(k+1)} = \mathcal{K}^{(k)} \cup \{p\}$.
6: **end for**
% The notation for indexing matrix columns is in Matlab style.

Assumption 1. *The matrix $\tilde{M} \in \mathbb{R}^{m\times n}$ is a near-separable matrix if $\tilde{M} = W[I_r \;\; H']\Pi + N$ with W full rank, I_r the r-by-r identity matrix, Π a permutation, $H'(:,i) \in \{x \in \mathbb{R}^r \mid x \geq 0, \sum_i x_i \leq 1\}$ for all i, and $\max_i ||N(:,i)||_2 \leq \epsilon$ for some $\epsilon > 0$.*

Theorem 1 ([10], Th. 3)**.** *Let $\tilde{M}$ satisfy Ass. 1. If $\epsilon \leq \mathcal{O}\left(\frac{\sigma_{\min}(W)}{\sqrt{r}\kappa^2(W)}\right)$, then SPA identifies the columns of W up to error $\mathcal{O}\left(\epsilon\,\kappa^2(W)\right)$, that is, the index set $\mathcal{K}$ identified by SPA satisfies $\max_j \min_{k\in\mathcal{K}} \left\|W(:,j) - \tilde{M}(:,k)\right\|_2 \leq \mathcal{O}\left(\epsilon\,\kappa^2(W)\right)$, where $\kappa(W) = \frac{\sigma_{\max}(W)}{\sigma_{\min}(W)}$ is the condition number of W.*

3 Preconditioning

Assume for simplicity that $m = r$. If we were given the full rank r-by-r matrix W, we could premultiply the input near-separable matrix $\tilde{M} = W[I_r, H'] + N$ with W^{-1} and obtain

$$\tilde{M}' = W^{-1}\tilde{M} = I_r[I_r, H'] + W^{-1}N.$$

The matrix $\tilde{M}'$ is also near-separable, and is now perfectly conditioned, although the noise $\epsilon' = \max_i ||W^{-1}N_{:i}||_2$ might have increased but by a factor of at most $\sigma^{-1}_{\min}(W)$. Under the same conditions as in Theorem 2, one can check that $\epsilon \leq \mathcal{O}\left(\frac{\sigma_{\min}(W)}{\sqrt{r}}\right)$ would imply that SPA identifies the columns of W up to error $\mathcal{O}\left(\epsilon\,\kappa(W)\right)$. This is a significant improvement compared to Theorem 2, both for the noise level ϵ (improvement of a factor $\kappa^{-2}(W)$) and the error (improvement of a factor $\kappa(W)$).

Of course, the matrix W is unknown, otherwise the problem would be solved. However, it can be shown that a minimum volume ellipsoid problem allows to approximate W^{-1} (up to orthogonal transformations): Given a matrix $\tilde{M} \in \mathbb{R}^{r\times n}$ of

rank r, we can formulate the minimum volume ellipsoid centered at the origin and containing the columns $\tilde{m}_i$ $1 \leq i \leq n$ of matrix $\tilde{M}$ as follows

$$A^* = \text{argmin}_{A \in \mathbb{S}^r_+} \log\det(A)^{-1} \text{ s.t. } \tilde{m}_i{}^T A \tilde{m}_i \leq 1 \quad \forall i. \tag{1}$$

This problem is SDP representable [11, p.222]. In [8], it is shown that, if $\tilde{M}$ satisfies Ass. 1, then $A^* \approx (WW^T)^{-1}$, hence factoring A^* (e.g., using the Cholesky decomposition) allows to recover W^{-1} approximately (up to orthogonal transformations). In fact, in the noiseless case (that is, for $N = 0$), $A^* = (WW^T)^{-1}$. Finally, our new proposed algorithm, referred to as preconditioned SPA (prec-SPA), works as follows: (i) Solve Problem (1) in order to approximately compute $P \approx W^{-1}$, (ii) Premultiply $\tilde{M}$ with the preconditioning to obtain $\tilde{M}' = P\tilde{M}$, and (iii) Apply SPA on the pre-conditioned matrix $\tilde{M}'$. We have the following robustness result:

Theorem 2 ([8], Th. 3)**.** *Let $\tilde{M}$ satisfy Ass. 1. If $\epsilon \leq \mathcal{O}\left(\frac{\sigma_{\min}(W)}{r\sqrt{r}}\right)$, then prec-SPA identifies a subset $\mathcal{K}$ so that $\tilde{M}(:,\mathcal{K})$ approximates the columns of W up to error $\mathcal{O}(\epsilon\,\kappa(W))$.*

4 Numerical Experiments

In this section, we illustrate the robustness of prec-SPA compared to the original SPA on some synthetic data sets. We also compare it to vertex component analysis (VCA) [12], a popular endmember extraction algorithm, and XRAY from [13]. The Matlab code is available at `https://sites.google.com/site/nicolasgillis/`. Please see [8] for more details and other numerical experiments.

The near-separable matrix $\tilde{M}$ is generated as follows: We take $m = r = 20$, and $n = 210$. The matrix W is generated using the `rand(.)` function of Matlab, that is, each entry is drawn uniformly at random in the interval $[0, 1]$. The matrix $H = [I_r, H']$ is such that H' has exactly two non-zero entries in each row equal to 0.5. Hence, the 190 data points are in the middle of two different columns of W. The noise is chosen such that the columns of W (that is, the first 20 columns of M) are not perturbed, while the 190 data points are moved towards the outside of the convex hull of the columns of W:

$$N = \delta\,[0_{20\times 20}, M(:,21) - \bar{w}, M(:,22) - \bar{w}, \ldots, M(:,n) - \bar{w}],$$

where $M = WH$ and $\bar{w} = \frac{1}{r}\sum_i w_i$ is the vertex centroid of the convex hull of the columns of W. These are the same near-separable matrices as in [10]. For different noise levels δ, we generate ten such matrices and Figure 1 reports the percentage of columns of W correctly identified by the different algorithms. Clearly, prec-SPA outperforms SPA, VCA and XRAY.

5 Conclusion

We have presented a way to precondition near-separable NMF matrices using semidefinite programming. This is turn allowed us to robustify near-separable NMF algorithms. In particular, the preconditioning makes the popular successive projection algorithm (SPA) provably more robust to noise, which we have illustrated on some synthetic data sets.

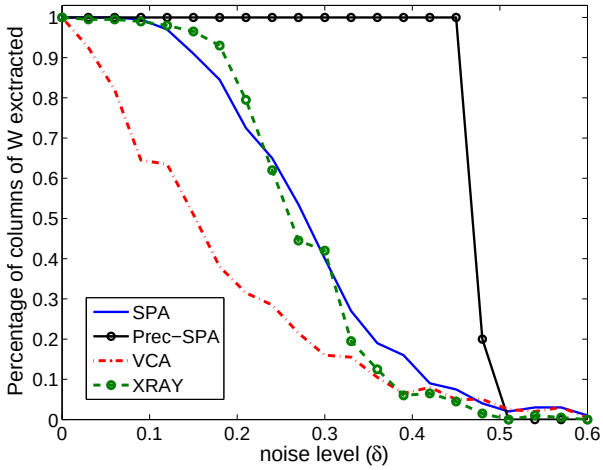

Figure 1: Comparison of the different near-separable NMF algorithms on synthetic data sets.

References

[1] S. Vavasis, "On the complexity of nonnegative matrix factorization," *SIAM J. on Optimization*, vol. 20, pp. 1364–1377, 2009.

[2] S. Arora, R. Ge, R. Kannan, and A. Moitra, "Computing a nonnegative matrix factorization – provably," in *STOC '12*, pp. 145–162.

[3] S. Arora, R. Ge, Y. Halpern, D. Mimno, A. Moitra, D. Sontag, Y. Wu, and M. Zhu, "A practical algorithm for topic modeling with provable guarantees," in *ICML '13*, vol. 28, pp. 280–288.

[4] W.-K. Ma, J. Bioucas-Dias, T.-H. Chan, N. Gillis, P. Gader, A. Plaza, A. Ambikapathi, and C.-Y. Chi, "A Signal Processing Perspective on Hyperspectral Unmixing: Insights from Remote Sensing," *IEEE Signal Process. Mag.*, vol. 31, pp. 67–81, 2014.

[5] T.-H. Chan, W.-K. Ma, C.-Y. Chi, and Y. Wang, "A convex analysis framework for blind separation of non-negative sources," *IEEE Trans. on Signal Processing*, vol. 56, pp. 5120–5134, 2008.

[6] E. Esser, M. Moller, S. Osher, G. Sapiro, and J. Xin, "A convex model for nonnegative matrix factorization and dimensionality reduction on physical space," *IEEE Trans. Image Process.*, vol. 21, pp. 3239–3252, 2012.

[7] E. Elhamifar, G. Sapiro, and R. Vidal, "See all by looking at a few: Sparse modeling for finding representative objects," in *CVPR '12*, pp. 1600–1607.

[8] N. Gillis and S. Vavasis, "Semidefinite programming based preconditioning for more robust near-separable nonnegative matrix factorization," 2013, arXiv:1310.2273.

[9] U. Araújo, B. Saldanha, R. Galvão, T. Yoneyama, H. Chame, and V. Visani, "The successive projections algorithm for variable selection in spectroscopic multicomponent analysis," *Chemom. and Intell. Lab. Syst.*, vol. 57, pp. 65–73, 2001.

[10] N. Gillis and S. Vavasis, "Fast and robust recursive algorithms for separable nonnegative matrix factorization," *IEEE Trans. Pattern Anal. Mach. Intell.*, 2013, doi:10.1109/TPAMI.2013.226.

[11] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge University Press, Cambridge, 2004.

[12] J. Nascimento and J. Bioucas-Dias, "Vertex component analysis: a fast algorithm to unmix hyperspectral data," *IEEE Trans. Geosci. Remote Sens.*, vol. 43, pp. 898–910, 2005.

[13] A. Kumar, V. Sindhwani, and P. Kambadur, "Fast conical hull algorithms for near-separable non-negative matrix factorization," in *ICML '13*, vol. 28, pp. 231–239.

Enhanced Recovery Conditions for OMP/OLS by Exploiting both Coherence and Decay

Cédric Herzet[1], Charles Soussen[2].

[1]INRIA Centre Rennes - Bretagne Atlantique, Rennes, France.

[2] Centre de Recherche en Automatique de Nancy, Université de Lorraine, Nancy, France.

1 Introduction

In this paper, we focus on two popular instances of greedy algorithms, namely orthogonal matching pursuit (OMP) [1] and orthogonal least squares[1] (OLS) [5, 6].

The suboptimal nature of OMP and OLS has led many researchers to study conditions under which these procedures succeed in recovering the true sparse vector. This question has been widely addressed for OMP in the recent years, including worst-case uniform [7, 8] and probabilistic analyses [9]. Although OLS has been known in the literature for a few decades (under different names [10]), exact recovery analysis for OLS have only appeared very recently, see [11, 12, 13].

Most of the existing works deal with *uniform* guarantees: these conditions ensure the success of OMP/OLS for a given sparsity level (or a given support) *irrespective* of the magnitude of the non-zero coefficients. In contrast with these works, we derive new guarantees of success accounting for the decay of the non-zero elements of the sparse vector. Our conditions are expressed in terms of the mutual coherence of the dictionary μ and encompass, as particular cases, some well-known results of the literature. The proofs of the results are reported in our technical report [14].

2 Context and Main Results

Let us assume that $\mathbf{y} \in \mathbb{R}^m$ is a (noisy) linear combination of k columns of $\mathbf{A} \in \mathbb{R}^{m\times n}$ indexed by $\mathcal{Q}^\star$, that is

$$\mathbf{y} = \mathbf{A}\mathbf{x} + \mathbf{w} \quad \text{with} \begin{cases} x_i \neq 0 \Leftrightarrow\ i \in \mathcal{Q}^\star \\ \mathrm{Card}\{\mathcal{Q}^\star\} = k \end{cases} \tag{1}$$

where $\mathbf{w} \in \mathbb{R}^m$ denotes some noise vector and $\mathrm{Card}\{\cdot\}$ stands for the cardinality operator. We assume that the columns $\mathbf{a}_i$ of the dictionary are normalized: $\|\mathbf{a}_i\|_2 = 1\ \forall i$.

Let us first consider the noiseless case ($\mathbf{w} = \mathbf{0}$). Theorem 1 provides sufficient conditions of success for OMP/OLS accounting for the possible decay of non-zero coefficients in $\mathbf{x}$. In our statement, we assume without loss of generality that

$$\mathcal{Q}^\star = \{1, 2, \ldots, k\}, \tag{2}$$

and

$$|x_1| \geq |x_2| \geq \ldots \geq |x_k| > 0. \tag{3}$$

Theorem 1. *If*

$$\mu < \frac{1}{k}, \tag{4}$$

and

$$|x_i| > \frac{2\mu(k-i)}{1-i\mu}|x_{i+1}| \qquad \forall i \in \{1, \ldots, k\}, \tag{5}$$

then Oxx selects atoms in $\mathcal{Q}^\star$ from noiseless data during the first k iterations.

In our technical report [14], we show that this kind of result can also be extended to the characterization of the success of OMP/OLS in the noisy setting ($\|\mathbf{w}\|_2 \leq \epsilon$) or the partial recovery of the support by the latter algorithms. In particular, we show that a well-known result by Donoho *et al.* [15, Th. 5.1] can be further relaxed as

Theorem 2. *If*

$$\mu < \frac{1}{2k-1}, \tag{6}$$

and

$$|x_i| > \frac{2\epsilon}{1-(2k-i)\mu} \qquad \forall i \in \{1, \ldots, k\}, \tag{7}$$

then Oxx selects atoms in $\mathcal{Q}^\star$ during the first k iterations.

More specifically, we see that the condition by Donoho *et al.* in [15, Th. 5.1] is sufficient for (7) to be satisfied and is therefore stronger than the conditions mentioned in Theorem 2. On the other hand, the two types of conditions become equivalent as soon as the nonzero elements of $\mathbf{x}$ have the same magnitude.

[1]The OLS algorithm is also known as forward selection [2], Order Recursive Matching Pursuit (ORMP) [3] and Optimized Orthogonal Matching Pursuit (OOMP) [4] in the literature.

References

[1] Y. C. Pati, R. Rezaiifar, and P. S. Krishnaprasad, "Orthogonal matching pursuit: recursive function approximation with applications to wavelet decomposition," in *Proc. 27th Ann. Asilomar Conf. Signals, Systems, and Computers*, 1993.

[2] Alan Miller, *Subset Selection in Regression, Second Editon*, Chapman and Hall/CRC, 2 edition, Apr. 2002.

[3] S. F. Cotter, J. Adler, B. D. Rao, and K. Kreutz-Delgado, "Forward sequential algorithms for best basis selection," *IEE Proc. Vision, Image and Signal Processing*, vol. 146, no. 5, pp. 235–244, Oct. 1999.

[4] L. Rebollo-Neira and D. Lowe, "Optimized orthogonal matching pursuit approach," *IEEE Signal Processing Letters*, vol. 9, no. 4, pp. 137–140, Apr. 2002.

[5] S. Chen, S. A. Billings, and W. Luo, "Orthogonal least squares methods and their application to non-linear system identification," *International Journal of Control*, vol. 50, no. 5, pp. 1873–1896, Nov. 1989.

[6] B. K. Natarajan, "Sparse approximate solutions to linear systems," *SIAM J. Comput.*, vol. 24, pp. 227–234, Apr. 1995.

[7] J. A. Tropp, "Greed is good: algorithmic results for sparse approximation," *IEEE Trans. Inf. Theory*, vol. 50, no. 10, pp. 2231–2242, Oct. 2004.

[8] M. A. Davenport and M. B. Wakin, "Analysis of orthogonal matching pursuit using the restricted isometry property," *IEEE Trans. Inf. Theory*, vol. 56, no. 9, pp. 4395–4401, 2010.

[9] J. A. Tropp and A. C. Gilbert, "Signal recovery from random measurements via orthogonal matching pursuit," *IEEE Trans. Inf. Theory*, vol. 53, no. 12, pp. 4655–4666, Dec. 2007.

[10] T. Blumensath and M. E. Davies, "On the difference between orthogonal matching pursuit and orthogonal least squares," Tech. Rep., University of Edinburgh, Mar. 2007.

[11] S. Foucart, "Stability and robustness of weak orthogonal matching pursuits," in *Recent advances in harmonic analysis and applications*, D. Bilyk, L. De Carli, A. Petukhov, A. M. Stokolos, and B. D. Wick, Eds. 2013, vol. 25, pp. 395–405, Springer proceedings in Mathematics & Statistics.

[12] C. Soussen, R. Gribonval, J. Idier, and C. Herzet, "Joint k-step analysis of orthogonal matching pursuit and orthogonal least squares," *IEEE Trans. Inf. Theory*, vol. 59, no. 5, pp. 3158–3174, May 2013.

[13] C. Herzet, C. Soussen, J. Idier, and R. Gribonval, "Exact recovery conditions for sparse representations with partial support information," *IEEE Trans. Inf. Theory*, vol. 59, no. 11, pp. 7509–7524, Nov. 2013.

[14] Cédric Herzet, Angélique Drémeau, and Charles Soussen, "Relaxed recovery conditions for OMP/OLS by exploiting both coherence and decay," Tech. Rep., available at http://arxiv.org/abs/1401.7533, Aug. 2014.

[15] D. L. Donoho, M. Elad, and V. N. Temlyakov, "Stable recovery of sparse overcomplete representations in the presence of noise," *IEEE Trans. Inf. Theory*, vol. 52, no. 1, pp. 6–18, Jan. 2006.

AUDIO DECLIPPING BY COSPARSE HARD THRESHOLDING

S. Kitić[†], *N. Bertin*[⋆], *R. Gribonval*[†]

[†]Inria , [⋆]IRISA - CNRS UMR 6074

ABSTRACT

Recovery of clipped audio signals is a very challenging inverse problem. Recently, it has been successfully addressed by several methods based on the sparse synthesis data model. In this work we propose an algorithm for enhancement of clipped audio signals that exploits the sparse analysis (*cosparse*) data model. Experiments on real audio data indicate that the algorithm has better signal restoration performance than state-of-the-art sparse synthesis declipping methods.

Index Terms— clipping, audio, sparse, cosparse, inverse problems

1. INTRODUCTION

Audio signals are prone to various degradations and clipping is among the most common ones. It may arise during recording due to the poor dynamic range of a microphone, or of analog-to-digital circuits, but it may be seen also as an extreme case of magnitude compression [1].

Clipping inevitably causes an information loss: in the time domain, the signal is "cut-off" when its amplitude is above some threshold. This manifests as an expansion (spread) in the frequency domain. An intuitive approach to de-clip clipped signals may be to "clean" their frequency spectrum and preserve only those frequencies needed to represent the original signal accurately. However, estimation of the correct support in the frequency domain is also a difficult problem. Hence, we need some additional prior information about the underlying signal.

We formalize clipping as the following element-wise operation on the original signal $\boldsymbol{x} \in \mathbb{R}^n$ (yielding the clipped signal $\bar{\boldsymbol{x}}$):

$$\bar{\boldsymbol{x}} = \operatorname{sign}(\boldsymbol{x}) \min(|\boldsymbol{x}|, \tau) \tag{1}$$

The threshold τ represents amplitude saturation level (here we assume *symmetric* clipping, *i.e.* the same threshold is applied to both positive and negative samples of $\boldsymbol{x}$).

This work was supported in part by the European Research Council, PLEASE project (ERC-StG-2011-277906).

[†]Centre Inria Rennes - Bretagne Atlantique, France

[⋆]Rennes, France

Let $\mathcal{M}_{\mathrm{r}}(\bar{\boldsymbol{x}})$ be the sampling operator which extracts unclipped, "*reliable*", samples, and $\mathcal{M}_{\mathrm{c}}^{+}(\bar{\boldsymbol{x}})$, $\mathcal{M}_{\mathrm{c}}^{-}(\bar{\boldsymbol{x}})$ be the operators which extract clipped samples with positive and negative amplitude, respectively. Knowing that (for the clipped positive/negative samples) the amplitude of the original signal has to be above (respectively, below) the threshold, we can formulate the inverse problem of signal de-clipping in the following way:

$$\begin{aligned} \textit{Find } \hat{\boldsymbol{x}} \in \mathbb{R}^n \textit{ such that: } & \mathcal{M}_{\mathrm{r}}(\hat{\boldsymbol{x}}) = \mathcal{M}_{\mathrm{r}}(\bar{\boldsymbol{x}}) \\ & \mathcal{M}_{\mathrm{c}}^{+}(\hat{\boldsymbol{x}}) \geq \mathcal{M}_{\mathrm{c}}^{+}(\bar{\boldsymbol{x}}) \\ & \mathcal{M}_{\mathrm{c}}^{-}(\hat{\boldsymbol{x}}) \leq \mathcal{M}_{\mathrm{c}}^{-}(\bar{\boldsymbol{x}}) \end{aligned} \tag{2}$$

This is an ill-posed inverse problem, since there are infinitely many possible solutions for $\hat{\boldsymbol{x}}$. We chose to regularize the problem by introducing a *sparsity hypothesis*.

2. SPARSE ANALYSIS DATA MODEL

The premise of sparse regularization is that many signals have low-dimensional representation, but not necessarily in their original domain. This assumption is sometimes sufficient to mitigate ill-posedness of the initial problem.

The *sparse analysis*, also known as the *cosparse* data model assumes that the signal in question can be "sparsified" by applying an adequate linear transform. Let $\boldsymbol{\Omega} \in \mathbb{R}^{p \times n}, (p \geq n)$ denote the matrix form of this sparsifying transform and let the linear constraints in (2) be denoted by Γ. We introduce the following optimization problem:

$$\text{minimize}_{\hat{\boldsymbol{x}}} \|\boldsymbol{\Omega}\hat{\boldsymbol{x}}\|_0 \text{ s.t. } \hat{\boldsymbol{x}} \text{ respects } \Gamma \tag{3}$$

However, minimizing this ℓ_0-"norm" is NP-hard, but there are convex relaxation and greedy algorithms [2] which may be used to approximate its solution.

The cosparse data model has only recently attracted the attention of scientific community, and to our best knowledge, it has not yet been exploited to solve the declipping problem.

3. COSPARSE DECLIPPING

Our approach is based on transposing the *Consistent Iterative Hard Thresholding* (Consistent IHT) algorithm [3] to the cosparse data model. Consistent IHT relies on applying hard

thresholding operator $H_K(\boldsymbol{s})$ [1] to the transform coefficient vector $\boldsymbol{s}$ in order to approximate sparse solutions. The transform coefficients are estimated by performing a gradient descent step of an objective function that incorporates the constraints expressed in Γ.

A pragmatical issue is that computing the k-cosparse projection is proven to be NP-hard [4]. To circumvent this problem, we take a different approach, largely based on the *Alternating Direction Method of Multipliers* [5] framework. The resulting algorithm is presented as pseudocode in algorithm 1 and will be referred to as Cosparse Declipping by Hard Thresholding (CoDec-HT).

Algorithm 1 COsparse DEClipping by Hard Thresholding

Require: $\boldsymbol{\Omega}, \bar{\boldsymbol{x}}, \Gamma, \varepsilon$
1: $\hat{\boldsymbol{x}}^{(0)} = \bar{\boldsymbol{x}}, \boldsymbol{u}^{(0)} = \boldsymbol{0}, k = 1$
2: $\hat{\boldsymbol{z}}^{(k)} = H_k\left(\boldsymbol{\Omega}\hat{\boldsymbol{x}}^{(k-1)} + \boldsymbol{u}^{(k-1)}\right)$
3: $\hat{\boldsymbol{x}}^{(k)} = \arg\min_{\tilde{\boldsymbol{x}}} \|\boldsymbol{\Omega}\tilde{\boldsymbol{x}} - \hat{\boldsymbol{z}}^{(k)} + \boldsymbol{u}^{(k-1)}\|_2^2$ s.t. $\tilde{\boldsymbol{x}}$ respects Γ
4: **if** $\|\boldsymbol{\Omega}\hat{\boldsymbol{x}}^{(k)} - \hat{\boldsymbol{z}}^{(k)} + \boldsymbol{u}^{(k-1)}\|_\infty \leq \varepsilon$ **then**
5: terminate
6: **else**
7: $\boldsymbol{u}^{(k)} = \boldsymbol{u}^{(k-1)} + \boldsymbol{\Omega}\hat{\boldsymbol{x}}^{(k)} - \hat{\boldsymbol{z}}^{(k)}$
8: $k \leftarrow k + 1$
9: go to 2
10: **end if**
11: **return** $\hat{\boldsymbol{x}}^{(k)}$

Mimicking the heuristics used in Consistent IHT, we relax the sparsity k of $\boldsymbol{z}$ through iterations, which allows to "learn" the unknown sparsity level of a given signal. Through iterations, the vectors $\boldsymbol{\Omega}\hat{\boldsymbol{x}}$ and $\hat{\boldsymbol{z}}$ eventually get close to each other, thus the signal estimate $\hat{\boldsymbol{x}}$ becomes approximately cosparse.

The analysis operator $\boldsymbol{\Omega}$ can be chosen as any (possibly overcomplete) transform which sparsifies the audio signal. In our experiments we used concatenated Discrete Cosine Transform and Discrete Sine Transform matrices.

4. EXPERIMENTS

We carried declipping experiments with 10-second excerpts from two wideband music pieces, sampled at 16kHz with 16bit encoding. Audio 1 is a rock piece with vocals, drums, two electric guitars and bass, while Audio 2 is an instrumental piece performed on an analogue synthesizer. The clipped versions of these signals are divided into overlapping frames of length 1024 before processing. Afterwards, the outputs are re-synthesized by the overlap and add scheme.

Beforehand, the audio is clipped given a predefined SNR_{inp} (signal-to-noise ratio of the input) value. The restoration performance is measured as the output signal-to-noise ratio (SNR_{out}):

$$\text{SNR}_{\text{inp}} = 20\log_{10}\frac{\|\boldsymbol{x}\|_2}{\|\boldsymbol{x}-\bar{\boldsymbol{x}}\|_2}, \text{SNR}_{\text{out}} = 20\log_{10}\frac{\|\boldsymbol{x}\|_2}{\|\boldsymbol{x}-\hat{\boldsymbol{x}}\|_2} \quad (4)$$

The proposed algorithm is compared against [3] and [6][2], which serve as a reference. Consistent IHT (with dictionary $\boldsymbol{\Psi} = \boldsymbol{\Omega}^H$) and the proposed algorithm are based on a similar algorithmic methodology and differ only by the underlying data model. Algorithm [6] exploits a structured prior termed *social sparsity* [7].

The results[3] presented in Figure 1 suggest that cosparsity-based hard thresholding outperforms synthesis-based one at all clipping levels. Interestingly, the SNR improvement rate of CoDec HT is roughly of the same order as for the social sparsity based algorithm, despite the fact that the latter uses a stronger prior.

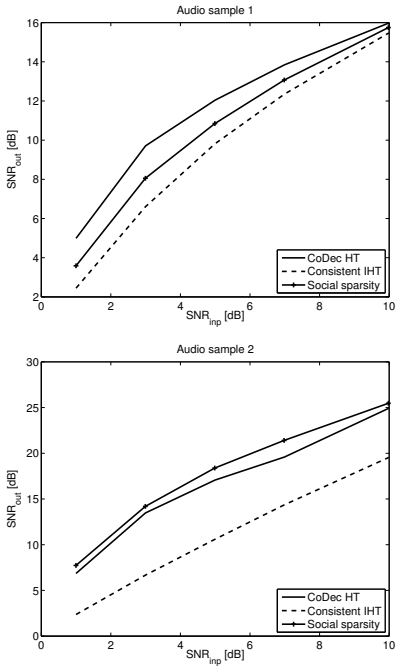

Fig. 1: Declipping results for two wideband audio tracks

5. CONCLUSION

Experimental results indicate that sparse analysis regularization can be successfully applied to the audio declipping inverse problem. Future work will aim at incorporating structured cosparsity in the data model to improve declipping performance, and at addressing extended magnitude corruption scenarios, such as dynamic range compression or soft clipping.

[1] $H_K(\boldsymbol{s})$ keeps K highest in absolute magnitude elements of $\boldsymbol{s}$ and sets the remaining to zero.

[2] We acknowledge M.Kowalski who provided his original code.

[3] Publicly available at <http://people.rennes.inria.fr/Srdan.Kitic/>.

6. REFERENCES

[1] J. O. Smith, *Physical audio signal processing: For virtual musical instruments and audio effects*, W3K Publishing, 2010.

[2] S. Nam, M. E. Davies, M. Elad, and R. Gribonval, "The Cosparse Analysis Model and Algorithms," *Applied and Computational Harmonic Analysis*, vol. 34, no. 1, pp. 30–56, 2013.

[3] S. Kitić, L. Jacques, N. Madhu, M. P. Hopwood, A. Spriet, and C. De Vleeschouwer, "Consistent iterative hard thresholding for signal declipping," in *Acoustics, Speech and Signal Processing (ICASSP), 2013 IEEE International Conference on*, 2013.

[4] A. M. Tillmann, R. Gribonval, and M. E. Pfetsch, "Projection onto the k-cosparse set is NP-hard," *CoRR*, vol. abs/1303.5305, 2013.

[5] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein, "Distributed optimization and statistical learning via the alternating direction method of multipliers," *Foundations and Trends® in Machine Learning*, vol. 3, no. 1, pp. 1–122, 2011.

[6] K. Siedenburg, M. Kowalski, and M. Dörfler, "Audio declipping with social sparsity," *Submitted to ICASSP, 2014*, 2013.

[7] M. Kowalski, K. Siedenburg, and M. Dörfler, "Social sparsity! neighborhood systems enrich structured shrinkage operators," 2012.

Strategies to learn computationally efficient and compact dictionaries

Luc Le Magoarou[1] and Rémi Gribonval[1]

[1] INRIA Rennes-Bretagne Atlantique, Campus de Beaulieu, 35042 Rennes, France.

***Abstract*— Dictionary learning is a branch of signal processing and machine learning that aims at expressing some given training data matrix as the multiplication of two factors: one dense matrix called dictionary and one sparse matrix being the representation of the data in the dictionary. The sparser the representation, the better the dictionary. However, manipulating the dictionary as a dense matrix can be computationally costly both in the learning process and later in the usage of this dictionary, thus limiting dictionary learning to relatively small-scale problems. In this paper we consider a general structure of dictionary allowing faster manipulation, and give an algorithm to learn such dictionaries over training data, as well as preliminary results showing the interest of our approach.**

1 Introduction

The use of a dictionary to sparsely represent a certain type of data goes back to almost two centuries with the Fourier transform that was designed to sparsely represent heat flow at that time [6]. The Fourier and the Hadamard transforms, as well as the wavelets to cite just a few rely on a "simple" mathematical formula that has been shown to yield fast algorithms [5, 9].

On the other hand, more recently have arisen algorithms that learn a dictionary directly over training data, without using an analytical formula (see [11] and references therein for a survey on the topic). They consider some data matrix $\mathbf{X} \in \mathbb{R}^{d\times n}$, which is the collection of n training vectors $\mathbf{x}_i \in \mathbb{R}^d$, and can be approximated sparsely in a certain dictionary:

$$\mathbf{X} \approx \mathbf{D}\mathbf{\Gamma}, \tag{1}$$

$\mathbf{\Gamma}$ having sparse columns. Such algorithms provide dictionaries that are well adapted to the data but they do not lead to fast algorithms because of the lack of structure of the dictionary, and storing $\mathbf{D}$ as a dense matrix may be impractical.

Bridging the gap between analytical dictionaries that are not necessarily well adapted to the data at hand but lead to fast algorithms, and learnt dictionaries that are very well adapted to the data but do not lead to fast algorithms is an important challenge. The dictionary learning community started recently to tackle the problem, mainly in [12] and [4], where the authors introduced new dictionary structures. The model we introduce below generalizes these works, as we will see in Sec.2.

In this paper we build on the very simple observation that the fast algorithms used to apply analytical transforms like the FFT or the DWT can be seen as consecutive multiplications of the input vector by sparse matrices. This fact implies that the analytical dictionaries associated with such fast transforms can be expressed as a product of sparse matrices, that is[1]:

$$\mathbf{D} = \prod_{j=1}^{M} \mathbf{S}_j. \tag{2}$$

This structure is precisely what makes possible the appearance of fast algorithms (such as the butterfly FFT for example). Knowing that, it makes sense to learn a dictionary that has this form and is the product of several sparse matrices, making it intrinsically fast and easier to store. The problem at hand is unfortunately highly non-convex and the sparsity constraint is non-smooth, but recent advances in optimization like the algorithm proposed in [2] can be adapted to such problems with convergence guarantees as we will show in the next sections.

[1]The product being taken from left to right: $\prod_{i=1}^{N} \mathbf{A}_i = \mathbf{A}_1 \cdots \mathbf{A}_N$

2 Model and optimization problem

Notation Throughout this paper, matrices are denoted by bold uppercase letters: $\mathbf{A}$. Vectors are denoted by bold lowercase letters: $\mathbf{a}$. The ith column of a matrix $\mathbf{A}$ is denoted by: $\mathbf{a}_i$. Sets are denoted by calligraphical symbols: $\mathcal{A}$.

As stated in the introduction, we propose a new dictionary structure that intrinsically leads to fast algorithms although the dictionary is learnt over some training data. Our goal is to find a dictionary that sparsely represent the data, being itself a multiplication of sparse matrices.

2.1 The matrix factorization point of view

Let $\mathbf{D} \in \mathbb{R}^{d\times a}$ be our dictionary with a atoms and $\mathbf{\Gamma} \in \mathbb{R}^{a\times n}$ the corresponding sparse representation matrix such that $\mathbf{X} \approx \mathbf{D}\mathbf{\Gamma}$. In order to meet the requirements and be intrinsically fast, $\mathbf{D}$ must take the form of equation (2), where the $\mathbf{S}_j$s are sparse matrices in $\mathbb{R}^{a_j\times a_{j+1}}$ with $a_1 = d$ and $a_{M+1} = a$. Now if we say that $\mathbf{\Gamma}$ is now called $\mathbf{S}_{M+1}$ for ease of notation, our goal is to find the sparse factors $\mathbf{S}_j$s such that:

$$\mathbf{X} \approx \prod_{j=1}^{M+1} \mathbf{S}_j. \tag{3}$$

We see with this equation that our structured dictionary learning problem amounts to a factorization of the data matrix into $M+1$ sparse factors. Actually, this model is quite general and can encompass those introduced in [12] and [4] as we will see. Such a multifactor representation of a data matrix has been introduced in the NMF framework in [8], for deep learning in [1, 10] and statistics in [7, 3].

2.2 Optimization objective

The proposed structured dictionary learning problem can be expressed as a constrained optimization problem. In its most general form, it can be stated as follows:

$$\begin{aligned} \underset{\mathbf{S}_j, \forall j\in\{1\cdots M+1\}}{\text{Minimize}} \quad & \left\|\mathbf{X} - \prod_{j=1}^{M+1} \mathbf{S}_j\right\|_F^2 \\ \text{Subject to} \quad & \mathbf{S}_j \in \mathcal{U}_j, \forall j \in \{1\cdots M+1\}, \end{aligned} \tag{4}$$

where the $\mathcal{U}_j$s are the sets in which each factor should lie.

In [12], the authors propose to constrain each atom of the dictionary to be a sparse linear combination of the atoms of a

so-called *base dictionary*. This base dictionary is assumed to be asociated with a fast algorithm (it takes the form of equation (2)), thus leading to fast manipulation of the whole learnt dictionary. It is actually equivalent to solving the problem (4) with the $M-1$ leftmost factors being known (the base dictionary), $\mathbf{S}_M$ and $\mathbf{S}_{M+1} = \mathbf{\Gamma}$ to estimate, and $\mathcal{U}_M$ and $\mathcal{U}_{M+1}$ denoting sparsity constraints. They give an algorithm to estimate jointly the dictionary and the representation. However, such a structure constrains the dictionary to be close to the base dictionary and thus does not provide full adaptability.

In [4], the authors constrain each atom in the dictionary to be the composition of several circular convolutions with sparse kernels, thus leading again to fast manipulation of the learnt dictionary. They give an algorithm to estimate the dictionary knowing the representation and the support of each kernel. It is actually equivalent to solving the problem (4) with $\mathbf{S}_{M+1}$ being known, $\mathcal{U}_i, \forall i \in \{1 \cdots M\}$ denoting the intersection of sparsity constraints with the set of circulant matrices, and the support of each factor being known.

We propose to handle problem (4) in a more general form, namely we will not assume that any factor is known, and we will consider any $\mathcal{U}_i$s provided that we are able to perform the projection onto it (note that in practice, only the $\mathcal{U}_i$s that are included in the set of sparse matrices will be of interest).

2.3 Algorithm

We propose to apply an algorithm stemming from recent advances in non-convex and non-smooth optimization to problem (4). The algorithm is introduced in [2] and called PALM (Proximal Alternating Linearized Minimization). PALM is made to handle objective functions that take the form:

$$\Psi(\mathbf{x}_1, \ldots, \mathbf{x}_p) := \sum_{i=1}^{p} f_i(\mathbf{x}_i) + H(\mathbf{x}_1, \ldots, \mathbf{x}_p), \qquad (5)$$

where each $\mathbf{x}_i \in \mathbb{R}^{n_i}$ is a block of variables, H is only assumed to be C^1 and the f_is are only assumed to be proper and lower-semicontinuous (possibly indicator functions of constraint sets). The idea behind PALM is to alternate updates between different blocks of variables (Gauss-Seidel), performing a proximal gradient step for each block. PALM is guaranteed to converge to a critical point of the objective function.

Problem (4) is obviously non-convex and thus the algorithm introduced in this paper is not guaranteed to converge toward the global minimum. Adapting PALM to the structured dictionary learning problem we just introduced is quite straightforward. One option is to take each $\mathbf{S}_j$ as a block of variables, so that H is the objective function of problem (4) and the f_is are the indicator functions of the sets $\mathcal{U}_j$. The proximal operator of the indicator function of a set $\mathcal{X}$ reducing to the projection operator onto this set $P_{\mathcal{X}}(.)$, the adaptation of PALM to our problem is given in Algorithm 1, where c_j^i is a stepsize depending on the Lipschitz modulus of the gradient $\nabla_{\mathbf{S}_j} H$ (see [2] for more details).

Algorithm 1 PALM for structured dictionary learning

for $i \in \{1 \cdots Niter\}$ **do**
  for $j \in \{1 \cdots M+1\}$ **do**
    Set $\mathbf{S}_j^{i+1} = P_{\mathcal{U}_j}\left(\mathbf{S}_j^i - \frac{1}{c_j^i} \nabla_{\mathbf{S}_j} H\left(\mathbf{S}_1^{i+1} \ldots \mathbf{S}_j^i \ldots \mathbf{S}_{M+1}^i\right)\right)$
  end for
end for

3 One practical example

Choosing the sets of constraints $\mathcal{U}_j$ is crucial in order to avoid as much as possible local minima, but in order to have a fast dictionary, they have to be subsets of the following (set of k_j-sparse matrices in $\mathbb{R}^{a_j \times a_{j+1}}$):

$$\mathcal{S}_{k_j}^{a_j \times a_{j+1}} = \{\mathbf{U} \in \mathbb{R}^{a_j \times a_{j+1}} : \sum_{i=1}^{a_{j+1}} \|\mathbf{u}_i\|_0^0 \leq k_j\}. \qquad (6)$$

In further work, we intend to set constraints such as the combination of sparsity and orthogonality, or some sort of structured sparsity, because we observed that the factorizations of the most famous analytical dictionaries are of this type.

However in this paper we will consider the simplest constraints, namely we just want the matrices $\mathbf{S}_j$s to have at most k_j nonzero entries. This means that we will have $\mathcal{U}_j = \mathcal{S}_{k_j}^{a_j \times a_{j+1}}$. In this simple configuration, the projection operator $P_{\mathcal{U}_j}(.)$ is simply the hard thresholding operator that sets all but the k_j greatest entries (in absolute value) to zero.

Complexity savings In the setting we just described, the multiplication of a vector by the dictionary or its adjoint (and its storage) would cost $\mathcal{O}(\sum_{j=1}^{M} k_j)$ operations instead of $\mathcal{O}(da)$ operations for a dictionary without structure. The interest here is twofold: first at the learning stage since this algorithm requires multiplications by the dictionary at each iteration, second at the usage stage since it involves again multiplications by the dictionary and its adjoint. It is thus necessary to tune M and the k_js in order to have a significant complexity gain, namely $\sum_{j=1}^{M} k_j \ll da$. For example in the FFT case with $d = a = 2^P$, $P \in \mathbb{N}$, the dictionary is factorized into $M = \log_2 d$ factors that are all $2d$-sparse, thus we perform multiplication in $\mathcal{O}(2d.\log_2 d)$ operations instead of $\mathcal{O}(d^2)$.

Preliminary result We are still testing different settings for the algorithm, but it shows promising results in tasks such as denoising. To illustrate this, we present an example of result with $d = 8$, $a = 20$, $n = 100$, $M = 3$, $k_4 = 100$, $k_1 = k_2 = k_3 = 20$ and $Niter$ = 1000. We generated randomly some data matrix $\mathbf{X}_0$ following the model of equation (3) and we added white Gaussian noise $\mathbf{B}$ with a signal-to-noise ratio of 6dB to obtain $\mathbf{X} = \mathbf{X}_0 + \mathbf{B}$. We ran Algorithm 1 in order to obtain $\hat{\mathbf{X}}_0$. We repeated the experiment 100 times and got an output signal-to-noise ratio of 7.7dB in average. We also made the interesting though expected observation that in approximation task (when we just want to factorize some matrix in sparse factors), there is a trade-off between sparsity of the factors and expressiveness of their product: the more non-zero entries we allow, the better the performances are, but the complexity is also increased.

4 Conclusion

In this abstract, we first presented the dictionary learning problem and the growing concern in the community to make it more computationally efficient. We then introduced a new model that is a generalization of two previously existing ones and that leads to intrinsically fast dictionaries. We presented an algorithm with convergence guarantees to a stationary point that is able to learn such type of dictionary over some training data. We finished by showing briefly some preliminary results of our approach. In the future, we intend to set new configurations for this method. We could for example modify the constraints in order to avoid as much as possible the numerous local minima inherent to the problem.

References

[1] S. Arora, A. Bhaskara, R. Ge, and T. Ma. Provable bounds for learning some deep representations. *CoRR*, abs/1310.6343, 2013.

[2] J. Bolte, S. Sabach, and M. Teboulle. Proximal alternating linearized minimization for nonconvex and nonsmooth problems. *Mathematical Programming*, pages 1–36, 2013.

[3] G. Cao, L. Bachega, and C. Bouman. The sparse matrix transform for covariance estimation and analysis of high dimensional signals. *Image Processing, IEEE Transactions on*, 20(3):625–640, 2011.

[4] O. Chabiron, F. Malgouyres, J.-Y. Tourneret, and N. Dobigeon. Toward Fast Transform Learning. Technical report, Nov. 2013.

[5] J. Cooley and J. Tukey. An algorithm for the machine calculation of complex fourier series. *Mathematics of Computation*, 19(90):297–301, 1965.

[6] J. Fourier. *Théorie analytique de la chaleur*. 1822.

[7] A. B. Lee, B. Nadler, and L. Wasserman. Treelets - an adaptive multi-scale basis for sparse unordered data. *The Annals of Applied Statistics*, 2(2):435–471, July 2008.

[8] S. Lyu and X. Wang. On algorithms for sparse multi-factor nmf. In *Advances in Neural Information Processing Systems 26*, pages 602–610. 2013.

[9] S. Mallat. A theory for multiresolution signal decomposition : the wavelet representation. *IEEE Transaction on Pattern Analysis and Machine Intelligence*, 11:674–693, June 1989.

[10] B. Neyshabur and R. Panigrahy. Sparse matrix factorization. *CoRR*, abs/1311.3315, 2013.

[11] R. Rubinstein, A. Bruckstein, and M. Elad. Dictionaries for Sparse Representation Modeling. *Proceedings of the IEEE*, 98(6):1045 –1057, 2010.

[12] R. Rubinstein, M. Zibulevsky, and M. Elad. Double sparsity: learning sparse dictionaries for sparse signal approximation. *IEEE Transactions on Signal Processing*, 58(3):1553–1564, Mar. 2010.

ITERATION-COMPLEXITY FOR INEXACT PROXIMAL SPLITTING ALGORITHMS

Jingwei Liang[1], *Jalal M. Fadili*[1] *and Gabriel Peyré*[2]

[1] GREYC, CNRS-ENSICAEN [2] CEREMADE, CNRS-Paris-Dauphine

ABSTRACT

In this work, we establish the iteration-complexity bounds (pointwise and ergodic) for inexact relaxed fixed point iteration built from nonexpansive operators, and then apply them to analyze the convergence rate of various proximal splitting methods in the literature. These include the forward-backward, Douglas-Rachford, ADMM and some primal-dual splitting methods. In particular, for the generalized forward-backward splitting algorithm [1], which covers the forward-backward and Douglas-Rachford algorithms as special cases, we develop an easily verifiable termination criterion for finding an approximate solution, which is a generalization of the termination criterion for the classical gradient descent method. We illustrate the usefulness of the above results by applying them to a large class of problems in signal and image processing.

Index Terms— Convex optimization, Proximal splitting, Convergence rates, Inverse problems.

1. INTRODUCTION

Many convex optimization problems in image and signal processing can be solved using a relaxed and inexact iterative scheme of the form

$$x^{k+1} = x^k + \lambda_k(Tx^k + \varepsilon^k - x^k), \tag{1.1}$$

where $T : \mathcal{H} \to \mathcal{H}$ is a nonexpansive operator on a Hilbert space $\mathcal{H}$, $\{\lambda_k\}_{k\in\mathbb{N}} \in]0, 1]$, and for $x^k \in \mathcal{H}$, $\varepsilon^k \in \mathcal{H}$ is the error when computing Tx^k.

One simple instance of such an algorithm is that of relaxed inexact gradient-descent for solving

$$\min_{x\in\mathcal{H}} f(x),$$

where f is proper convex and has β^{-1}-Lipschitz continuous gradient, in which case $T = \mathrm{Id} - \gamma\nabla f$, for $\gamma \in]0, 2\beta[$. In this case, the error ε^k is that when evaluating ∇f at x^k.

Many structured convex optimization problems boil down to implementing the same type of iterations. Consider for instance

$$\min_{x\in\mathcal{H}} \big\{\Phi(x) = f(x) + \textstyle\sum_{i=1}^n h_i(x)\big\}, \tag{1.2}$$

where $f \in \Gamma_0(\mathcal{H})$ has β^{-1}-Lipschitz continuous gradient, $h_i \in \Gamma_0(\mathcal{H})$, and $\Gamma_0(\mathcal{H})$ is the class of lower semicontinuous, proper, convex functions from a Hilbert space $\mathcal{H}$ to $]-\infty, +\infty]$. We assume that some technical domain qualification conditions are verified. Problem (1.2) was considered in [1].

In the last decades, based on the notion of the proximity operator [2], and assuming that the functions h_i are simple (*i.e.*their proximity operators are easily computable), a wide range of proximal splitting algorithms have been proposed to solve (1.2). One can cite for instance the forward-backward splitting method (FBS) valid for $n = 1$, the Douglas-Rachford splitting method (DRS) that applies for $f = 0$, a generalized Forward-Backward splitting method (GFB), or primal-dual splitting methods. See [3] for a comprehensive account.

In this paper, we first establish the iteration complexity of the inexact relaxed fixed point iteration (1.1). We then build upon this result to provide iteration complexity bounds for several proximal splitting algorithms, from which an easily verifiable termination criterion for finding an approximate solution will be provided.

2. ITERATION COMPLEXITY OF INEXACT NONEXPANSIVE OPERATOR

Define

$$e^k = (\mathrm{Id} - T)x^k. \tag{2.1}$$

Let $\tau_k = \lambda_k(1 - \lambda_k)$, $\underline{\tau} = \inf_{k\in\mathbb{N}} \tau_k$, $\overline{\tau} = \sup_{k\in\mathbb{N}} \tau_k$, and $\nu_1 = 2\sup_{k\in\mathbb{N}} \|T_k x^k - x^\star\| + \sup_{k\in\mathbb{N}} \lambda_k\|\varepsilon^k\|$, $\nu_2 = 2\sup_{k\in\mathbb{N}} \|e^k - e^{k+1}\|$, where $x^\star \in \mathrm{fix}T$ and $\mathrm{fix}T$ is the set of fixed points of T.

Theorem 2.1 (Pointwise iteration complexity bound of (1.1)). *Let d_0 be the distance from the starting point x^0 to the solution set $\mathrm{fix}T$, if the following conditions,*

(a) $\mathrm{fix}T \neq \varnothing$,

(b) $\lambda_k \in]0, 1[$, *such that* $\underline{\tau} > 0$ *and* $\sum_{k\in\mathbb{N}} \tau_k = +\infty$,

(c) $\sum_{k\in\mathbb{N}}(k+1)\|\varepsilon^k\| < +\infty$,

are satisfied, then

This work has been supported by the ERC project SIGMA-Vision and l'Institut Universitaire de France.

(i) $\{e^k\}_{k\in\mathbb{N}}$ *converges strongly to* 0, $\{x^k\}_{k\in\mathbb{N}}$ *converges weakly to a point* $x^\star \in \mathrm{fix}T$*;*

(ii) *Denote* $C_1 = \nu_1 \sum_{j\in\mathbb{N}} \lambda_j \|\varepsilon^j\| + \nu_2 \overline{\tau} \sum_{\ell\in\mathbb{N}} (\ell+1)\|\varepsilon^\ell\| < +\infty$*, then*

$$\|e^k\| \leqslant \sqrt{\frac{d_0^2 + C_1}{\underline{\tau}(k+1)}}. \tag{2.2}$$

Next we present the ergodic iteration complexity bound of (1.1), define $\Lambda_k = \sum_{j=0}^k \lambda_j$, and $\bar{e}^k = \frac{1}{\Lambda_k}\sum_{j=0}^k \lambda_j e^j$.

Theorem 2.2 (Ergodic iteration complexity bound of (1.1)). *Assume that condition Theorem 2.1 (c) holds and let* $C_3 = \sum_{j\in\mathbb{N}} \lambda_j \|\varepsilon^j\| < +\infty$*, then,*

$$\|\bar{e}^k\| \leqslant \frac{2d_0 + C_3}{\Lambda_k}. \tag{2.3}$$

A special class of nonexpansive operators is that of α-averaged operators [3] for $\alpha \in]0,1[$. The above two complexity bounds obviously apply, where condition Theorem 2.1 (b) can be changed to

$$\lambda_k \in]0, \frac{1}{\alpha}[\text{ s.t. } \inf_{k\in\mathbb{N}} \lambda_k(\frac{1}{\alpha} - \lambda_k) > 0, \ \sum_{k\in\mathbb{N}} \lambda_k(\frac{1}{\alpha} - \lambda_k) = +\infty \ .$$

3. APPLICATIONS TO PROXIMAL SPLITTING

FBS Suppose that $n = 1$ in (1.2). Then FBS with a fixed step-size corresponds to $T = \mathrm{prox}_{\gamma h_1} \circ (\mathrm{Id} - \gamma\nabla f)$, $\gamma \in]0, 2\beta[$, $\lambda_k \in]0, \frac{4\beta-\gamma}{2\beta}[$, and ε^k is the error when evaluating both the proximity operator and the gradient. In this case, setting $g^{k+1} = \frac{1}{\gamma}(x^k - x^{k+1}) - \nabla f(x^k)$, it follows that

$$g^{k+1} \in h_1(x^{k+1}), \text{ and } \|g^{k+1} + \nabla f(x^{k+1})\|^2 = \mathcal{O}(1/k) \ .$$

In plain words, this means that $\mathcal{O}(1/\epsilon)$ iterations are needed to find a pair $(x, g \in \partial h_1(x))$ with the termination criterion $\|g + \nabla f(x)\|^2 \leqslant \epsilon$. This iteration-complexity improves to $\mathcal{O}(1/\sqrt{\epsilon})$ in ergodic sense for the same termination criterion.

GFB Consider now $n > 1$ in (1.2). [1] proposed a generalization in a product space $\mathcal{H}^n$ of the FBS scheme to solve (1.2). Tn the GFB, T takes a more intricate form, omitted here for space limitation, and ε^k absorbs the error when computing both the proximity operators $\mathrm{prox}_{\gamma h_i}$ and the gradient. Let γ and λ_k be chosen as for FBS. We prove that $O(1/\epsilon)$ iterations are needed to find a pair $((u_i)_{1\leqslant i\leqslant n}, g)$ with the termination criterion $\|g + \nabla f(\sum_i \omega_i u_i)\|^2 \leqslant \epsilon$, where $g \in \sum_{i=1}^n \partial h_i(u_i)$, see [4] for details. Only $\mathcal{O}(1/\sqrt{\epsilon})$ are however needed in ergodic sense.

Other splitting schemes Similar complexity bounds were established for DRS, ADMM and some some primal-dual splitting methods to solve more general monotone inclusion problems that encompass (1.2) as a special case; see [5] for details.

4. NUMERICAL EXPERIMENTS

As an illustrative example, let's consider the principal component pursuit (PCP) problem [6], and apply it to decompose a video sequence into its background and foreground components. The rationale behind this is that since the background is virtually the same in all frames, if the latter are stacked as columns of a matrix, it is likely to be low-rank (even of rank 1 for perfectly constant background). On the other hand, moving objects appear occasionally on each frame and occupy only a small fraction of it. Thus the corresponding component would be sparse. The results are depicted on Fig. 1.

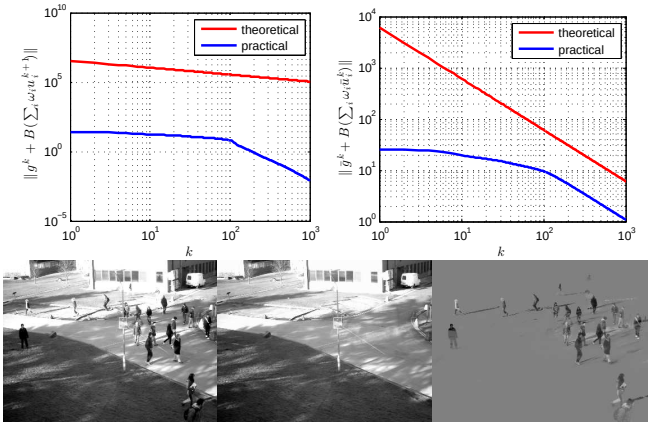

Fig. 1: Top: observed pointwise and ergodic rates and theoretical bounds for the GFB applied to the PCP problem. Bottom: original video frame out of 300 (left), recovered low-rank background (middle) and sparse foreground (right).

5. REFERENCES

[1] Hugo Raguet, Jalal M. Fadili, and Gabriel Peyré, "Generalized forward-backward splitting," *SIAM Journal on Imaging Sciences*, vol. 6, no. 3, pp. 1199â–1226, 2013.

[2] Jean-Jacques Moreau, "Proximité et dualité dans un espace hilbertien," *Bulletin de la Société mathématique de France*, vol. 93, pp. 273–299, 1965.

[3] Heinz H Bauschke and Patrick L Combettes, *Convex analysis and monotone operator theory in Hilbert spaces*, Springer, 2011.

[4] Jingwei Liang, Jalal M. Fadili, and Gabriel Peyré, "Iteration-complexity of a generalized forward backward splitting algorithm," *arXiv:1310.6636*, 2013.

[5] Jingwei Liang, Jalal M. Fadili, and Gabriel Peyré, "Iteration-complexity with inexact nonexpansive operators: Application to monotone inclusion," Tech. Rep., 2013.

[6] Emmanuel J Candès, Xiaodong Li, Yi Ma, and John Wright, "Robust principal component analysis?," *Journal of the ACM (JACM)*, vol. 58, no. 3, pp. 11, 2011.

CALIBRATION AND IMAGING THROUGH SCATTERING MEDIA

Antoine Liutkus *David Martina* *Sylvain Gigan* *Laurent Daudet*

Institut Langevin, ESPCI ParisTech, Paris Diderot Univ., UPMC Univ. Paris 6
CNRS UMR 7587, Paris, France

ABSTRACT

Multiply scattering media in optics such as thin layers of ZnO — essentially white paint — behave as natural random linear multiplexers that can be exploited to yield *opaque lenses* or compressive imagers that beat their coherent counterparts both in terms of performance and ease of implementation. As such, they appear as a very promising research topic in many domains of optics. Still, current research is hindered by the fact that the linear model relating the input and the output of a multiply scattering material only holds for the complex optical field, which is difficult to measure without an interferometric experimental setup. Indeed, in a practical setup, only its amplitude is easily captured by a CCD camera. In this study, we show how such intensity measurements and digital micromirror devices can be exploited both for calibration, arbitrary focusing and compressive imaging using scattering materials.

***Index Terms*—intensity-based calibration, focusing, compressive sensing, optical imaging, scattering media**

I. INTRODUCTION

Whereas optical scattering is usually seen as a time-varying nuisance, for instance when imaging through turbid media, the recent results of wave control in stable complex material have largely demonstrated that it could also be exploited, for example so as to build focusing systems that beat their coherent counterparts in terms of resolution [9]. In the context of compressive sensing for instance (CS, see [2]), such materials perform an efficient randomized multiplexing of the object into several sensors and were recently considered as analog randomizers [6]. Using such naturally occurring randomness instead of digitally-controlled pseudo-random sequences as in the single-pixel camera [3] permits both to drop the need for a designer to craft the randomizer himself, but also yields very compact sensing devices.

However, some practical difficulties come along using multiply scattering materials as multiplexing devices for focusing and imaging. Most importantly, the sampling mechanism is completely unknown *a priori* when using a multiply scattering material. Several studies have demonstrated that the input-output relationship of multiply scattering materials can very well be modelled through a linear Transmission Matrix (TM, see [8]), which can be used both for focusing [8], overdetermined imaging [7] and CS [6].

In this study, we go further previous work that only consider linear measurements and show how intensity-based calibration can be performed efficiently, which permits to drop the need for interferometric experimental setups and to simply use CCD cameras as sensors.

With the estimated TM, we show how focusing can be performed, where the controlling device is explicitly assumed to be a binary DMD. In comparison to the work presented in [8], [9], [1], the proposed approach has several advantages. First it does not require an interferometric mechanism nor an expensive Spatial Light Modulator (SLM) as in [8]. Second, it permits to build arbitrary output patterns to the system whereas [9], [1] only permit one single focusing point.

This work was supported by the European Research Council (Grant 278025), the Emergence(s) program from the City of Paris, and LABEX WIFI (Laboratory of Excellence within the French Program "Investments for the Future") under references ANR-10-LABX-24 and ANR-10-IDEX-0001-02 PSL*.

II. LINEAR MEASUREMENTS

II-A. Theoretical background

Due to space constraints, we do not detail here the theory underlying the present study. The interested reader is refferred to [6] for more details. In short, the input to the system is a vector $x \in \mathbb{R}^N$. The measurements are given as a vector $y \in \mathbb{R}^M$, given by:

$$y = Hx, \tag{1}$$

where the $M \times N$ matrix H is called the transmission matrix (TM) and accounts for the action of the multiply scattering material. In practice, the TM H has some interesting properties. In particular, its entries can be assumed Gaussian independent and identically distributed (i.i.d.) random variables (see [6] and references therein).

With H estimated, an arbitrary output pattern y may be obtained by settling the input to $x = H^\dagger y$, where $H^\dagger$ denotes the phase conjugate of H in a time-reversal interpretation, or any regularized inverse (see [6] for references). In practice, x cannot take any arbitrary values, but is often constrained to lie either on the complex unit circle (as in the case of a phase-only modulation using a SLM) or to be binary $x \in \{0,1\}^N$ as in the case of a binary DMD.

II-B. Calibration: measuring the measurement matrix

Assume that during a calibration stage, we can control the input x and measure the corresponding output y. Doing this L times and stacking the inputs and the outputs into matrices X and Y respectively, we have $Y = HX$. If we assume that these measurements are given with additive white Gaussian noise of variance σ^2, we can estimate H through $\hat{H} = YX^H \left(XX^H + \sigma^2 I\right)^{-1}$, where I is the $N \times N$ identity matrix. Correlated measurement noise can easily be taken into account, as well as a prior distribution over H in a Bayesian setting.

II-C. Experimental phase transitions under noise

All CS reconstruction algorithms exhibit some level of robustness to noise. However, when noise becomes prominent, performance of CS eventually drops. To overcome those limitations, it has been shown that it is advantageous to benefit from Multiple Measurement Vectors (MMV, [4]). In the MMV setting, $P > 1$ different measurements are available that correspond to the output of the system where the corresponding input all share the same support. In figure 1, we show observed performance of our compressive optical imager with linear measurements, using a varying number P of measurements per sparse object to reconstruct.

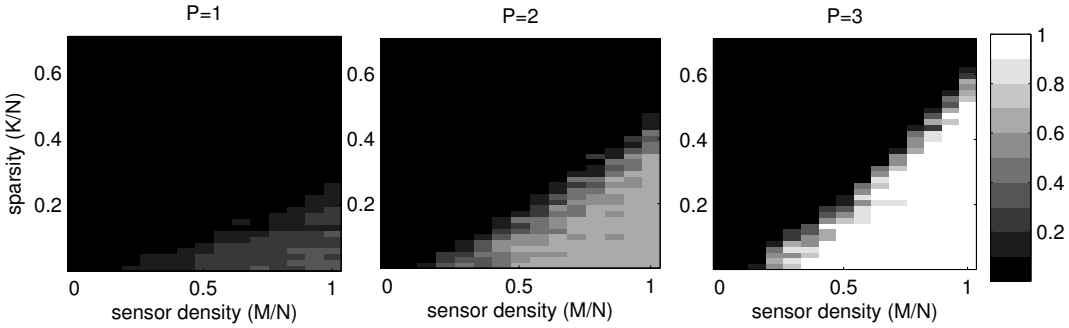

Fig. 1. Experimental performance of CS reconstruction with linear measurements in our real optical setup. Due to the presence of strong noise, only MMV at $P = 3$ permits reconstruction with high probability.

III. AMPLITUDE MEASUREMENTS

III-A. Theoretical background

In the case of amplitude measurements as performed by a CCD camera for instance, the observed output is $I = |y| = |Hx|$. In that case, the TM H cannot be estimated simply be a least-squared error procedure as above because of the non-linearity introduced by the amplitude operator. However, it can be estimated in an iterative fashion by a dedicated algorithm that we cannot detail here due to space constraints.

III-B. Focusing with a binary DMD

Once H has been estimated, it can be used for both imaging and focusing. In the case of focusing, we want a particular output y and need the corresponding input x. When the input is to be binary as in the case of a DMD, we propose to use an *antisparse coding* scheme [5] that looks for x by solving

$$\hat{x} = \underset{x}{\operatorname{argmin}} \left[\|Hx - y\|_2 + \|x\|_\infty \right]. \qquad (2)$$

Whereas the first term in 2 accounts for the focusing quality, the second term tends to minimize the maximal value of x,which in turns leads to x having only two values : $\pm \|x\|_\infty$, thus being binary.

In figure 2, we show an actual preliminary example of the focusing performance obtained with the proposed approach, where we want to focus on three target points at the same time. Intensity-based calibration was achieved with $4N$ measurements ($N = 900$ here). As can be seen, the proposed approach permits efficient focusing.

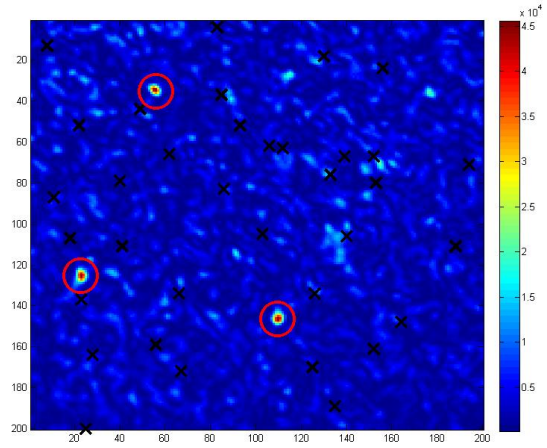

Fig. 2. Experimental focusing using a DMD and a scattering material, with intensity-based calibration. The red-circles are the target points, while black crosses are control points where intensity is minimized.

IV. CONCLUSION

In this paper, we have briefly summarized some important trends in optical imaging with multiply scattering materials. Those materials perform efficient multiplexing of input lights and are thus particularly efficient in the context of compressed sensing. We have shown how recent results and algorithms from the *sparsity framework* may be used to yield efficient optical devices that permit both focusing and imaging with a limited number of intensity calibration measurements.

V. REFERENCES

[1] D. Akbulut, T. Huisman, E. van Putten, W. Vos, and A. Mosk. Focusing light through random photonic media by binary amplitude modulation. *arXiv preprint arXiv:1101.2860*, 2011.

[2] D.L. Donoho. Compressed sensing. *Information Theory, IEEE Transactions on*, 52(4):1289–1306, 2006.

[3] M. Duarte, M. Davenport, D. Takhar, J. N Laska, T. Sun, K. F Kelly, and R. Baraniuk. Single-pixel imaging via compressive sampling. *Signal Processing Magazine, IEEE*, 25(2):83–91, 2008.

[4] Y. Eldar and H. Rauhut. Average case analysis of multichannel sparse recovery using convex relaxation. *Information Theory, IEEE Transactions on*, 56(1):505–519, 2010.

[5] H. Jégou, T. Furon, and J-J. Fuchs. Anti-sparse coding for approximate nearest neighbor search. In *Acoustics, Speech and Signal Processing (ICASSP), 2012 IEEE International Conference on*, pages 2029–2032. IEEE, 2012.

[6] A. Liutkus, D. Martina, S. Popoff, G. Chardon, O. Katz, G. Lerosey, S. Gigan, L. Daudet, and I. Carron. Imaging with nature: A universal analog compressive imager using a multiply scattering medium. *arXiv preprint arXiv:1309.0425*, 2013.

[7] S. Popoff, G. Lerosey, M. Fink, A. Boccara, and S. Gigan. Image transmission through an opaque material. *Nature Communications*, 1:81, 2010.

[8] S. Popoff, G. Lerosey, M. Fink, A. Boccara, and S. Gigan. Controlling light through optical disordered media: transmission matrix approach. *New Journal of Physics*, 13(12):123021, 2011.

[9] I. Vellekoop, A. Lagendijk, and A. Mosk. Exploiting disorder for perfect focusing. *Nature Photonics*, 4(5):320–322, 2010.

Multiscale Dictionary and Manifold Learning: Non-Asymptotic Bounds for the Geometric Multi-Resolution Analysis

Mauro Maggioni[1,2,3], Stanislav Minsker[1,4] and Nate Strawn[1,2].

[1]Department of Mathematics, [2]Electrical and Computer Engineering, [3]Computer Science, [4]Statistics, Duke University, Durham, N.C., U.S.A.

***Abstract*— Data in high-dimensional spaces often aggregate near inherently low-dimensional structures. Manifold learning and dictionary learning motivate numerous techniques for exploiting such low intrinsic dimensionality. In manifold learning, a low-dimensional manifold constitutes the underlying structure of the data, and one "learns" this manifold from samples. In dictionary learning, the goal is to construct a dictionary (i.e. a set of vectors) such that every data point may be expressed as a sparse (compressible) linear combination of elements from this set. Manifold learning and dictionary learning are generally considered as separate problems, however, Geometric Multi-Resolution Analysis (GMRA) [2] provides a method for approximating (in a robust multiscale fashion) a low-dimensional set around which data may concentrate, while simultaneously producing a dictionary for representing the data sparsely.**

In this work we introduce an estimator for low-dimensional sets supporting the data constructed from the GMRA approximations. We prove near-optimal finite sample bounds on its performance, and we also rigorously establish the robustness of this procedure with respect to both noise and model error. For data supported on a low-dimensional manifold, we demonstrate that this procedure admits error bounds primarily dependent upon the global parameters of the manifold: its intrinsic dimension and reach.

1 Introduction

In many high-dimensional data analysis problems, efficient representations dramatically boost the performance and efficiency of learning algorithms. In particular, it is often convenient to work with *sparse* representations of the data, meaning that there exists a dictionary (a large collection $\Phi = \{\phi_i\}_{i=1}^N$ of fixed vectors) such that every observation can be represented as a linear combination of the elements of Φ with few nonzero coefficients (e.g., see [1, 7, 12]). In practice, the dictionary Φ is not known, and must be inferred algorithmically from the training data $\mathcal{X} = \{X_1, \dots, X_n\}$ consisting of n i.i.d. samples from some underlying distribution Π. In most cases, this approach results in a random error due to the fact that the dictionary is learned based on the empirical approximation of Π.

When the *intrinsic dimension* of the analyzed data is small, sparse representations generally become available. One of the simplest geometric models is the assumption that the $\mathbb{R}^D$-valued observations lie close to a d-dimensional linear subspace $V \subset \mathbb{R}^D$, where $d \ll D$. Classical tools such as PCA [17, 21] can estimate V and use its basis as a dictionary. One may generalize this model by using $K \geq 2$ low-dimensional subspaces, and the resulting complexity of this generalization inspired intensive research over the past decade (e.g. [5, 15, 23, 25–27] and references therein). Another popular assumption is that the data is supported on a manifold $\mathcal{M}$ of dimension d isometrically embedded in $\mathbb{R}^D$ [3, 8–11, 13, 14, 18, 19, 22, 24, 28]. Under this structural assumption, one can (locally) represent the point on the manifold $\mathcal{M}$ in the basis of the tangent space which has d elements. The GMRA procedure of [2, 6] provides a multiscale, computationally efficient way to construct such bases, find the corresponding coordinates for a given point, and that paper demonstrates approximation guarantees in the infinite sample limit.

The main goal of this work is to obtain tight *non-asymptotic* bounds for the reconstruction error achieved by the GMRA-based method under the assumption that the distribution Π of the data possesses certain geometric properties. In particular, we verify the conditions for our guarantees for the case when Π is comparable with uniform distribution on a smooth submanifold of $\mathbb{R}^D$ (or a thin tube around a submanifold), and the error rate primarily depends on the intrinsic dimension of the manifold. Detailed discussions and proofs are provided in [20].

2 Definitions and Notation

The following notation is used throughout the paper: $\|\cdot\|_2$ denotes the standard Euclidean norm in $\mathbb{R}^D$, Proj_V stands for the orthogonal projection on a linear subspace $V \subseteq \mathbb{R}^D$, $\dim(V)$ for its dimension and $V^\perp$ for the orthogonal complement. Given a matrix $A \in \mathbb{R}^{D\times D}$, $\mathrm{tr}\,(A)$ denotes its trace, $\|\cdot\|$ – its operator norm and A^T – its transpose.

A GMRA with respect to the probability measure Π consists of a sequence of (nonlinear) operators $\{P_j : \mathbb{R}^D \to \mathbb{R}^D\}_{j\geq 0}$. For each "resolution level" $j \geq 0$, P_j is uniquely defined by a collection of pairs of subsets and affine projections, $\{(C_{j,k}, P_{j,k})\}_{k=1}^{N(j)}$, where the subsets $\{C_{j,k}\}_{k=1}^{N(j)}$ form a Borel-measurable disjoint partition of $\mathbb{R}^D$ and $P_{j,k}(x) := c_{j,k} + \mathrm{Proj}_{V_{j,k}}(x - c_{j,k})$, where $c_{j,k} \in \mathbb{R}^D$ and $V_{j,k}$ are defined as follows. Let $\mathbb{E}_{j,k}$ stand for the expectation with respect to the conditional distribution $d\Pi_{j,k}(x) = d\Pi(x|x \in C_{j,k})$. Then $c_{j,k} = \mathbb{E}_{j,k}[x]$ and

$$V_{j,k} = \underset{\dim(V)=d}{\mathrm{argmin}}\ \mathbb{E}_{j,k}\,\|x - c_{j,k} - \mathrm{Proj}_V(x - c_{j,k})\|_2^2\ .$$

In other words, $c_{j,k}$ is the conditional mean and $V_{j,k}$ is the subspace spanned by eigenvectors corresponding to d largest eigenvalues of the conditional covariance matrix

$$\Sigma_{j,k} = \mathbb{E}_{j,k}(x - c_{j,k})(x - c_{j,k})^T. \tag{2.1}$$

Given such a $\{(C_{j,k}, P_{j,k})\}_{k=1}^{N(j)}$, we define

$$P_j(x) := \sum_{k=1}^{N(j)} I\{x \in C_{j,k}\}P_{j,k}(x)$$

where $I\{x \in C_{j,k}\}$ is the indicator function of the set $C_{j,k}$.

It was shown in [2] that if Π is supported on a smooth, compact d-dimensional submanifold $\mathcal{M} \hookrightarrow \mathbb{R}^D$, and if the partitions $\{C_{j,k}\}_{j=1}^{N(j)}$ satisfy some regularity conditions for each j,

then, for any $x \in \mathcal{M}$, $\|x - P_j(x)\|_2 \leq C(\mathcal{M})2^{-2j}$ for all $j \geq j_0(\mathcal{M})$. This means that the operators P_j provide an efficient "compression scheme" for $x \in \mathcal{M}$, in the sense that every x can be well-approximated by a linear combination of at most $d+1$ vectors. In practice, Π is unknown and we only have access to the *training data* $\mathcal{X} = \{X_1, \ldots, X_n\}$, which are assumed to be i.i.d. with distribution Π. In this case, operators P_j are replaced by their estimators

$$\widehat{P}_j(x) := \sum_{k=1}^{N(j)} I\{x \in C_{j,k}\}\widehat{P}_{j,k}(x)$$

where $\{C_{j,k}\}_{k=1}^{N(j)}$ is a suitable partition of $\mathbb{R}^D$ obtained from the data, $\mathcal{X}_{j,k} = C_{j,k} \cap \mathcal{X}$, $|\mathcal{X}_{j,k}|$ is the cardinality of $\mathcal{X}_{j,k}$,

$$\begin{aligned}
&\widehat{c}_{j,k} = \frac{1}{|\mathcal{X}_{j,k}|}\sum_{x\in\mathcal{X}_{j,k}} x,\\
&\widehat{P}_{j,k}(x) := \widehat{c}_{j,k} + \mathrm{Proj}_{\widehat{V}_{j,k}}(x - \widehat{c}_{j,k}),\\
&\widehat{V}_{j,k} = \operatorname*{argmin}_{\dim(V)=d} \frac{1}{|\mathcal{X}_{j,k}|}\sum_{x\in\mathcal{X}_{j,k}} \|x - \widehat{c}_{j,k} - \mathrm{Proj}_V(x - \widehat{c}_{j,k})\|_2^2,
\end{aligned} \tag{2.2}$$

We call these $\widehat{P}_j$ the *empirical GMRA*.

3 Assumptions and Main Results

3.1 $L_2(\Pi)$ error bounds for GMRA

We now state the sufficient conditions on the distribution Π and the partition $\{C_{j,k}\}_{j=1}^{N(j)}$ which imply our $L_2(\Pi)$ error bound for the GMRA approximation. Suppose that for all $j_{\min} \leq j \leq j_{\max}$ the following hold:

(A1) There exists an integer $1 \leq d \leq D$ and a positive constant $\theta_1 = \theta_1(\Pi)$ such that for all $k = 1, \ldots, N(j)$,

$$\Pi(C_{j,k}) \geq \theta_1 2^{-jd}.$$

(A2) There is a positive constant $\theta_2 = \theta_3(\Pi)$ such that for all $k = 1 \ldots N(j)$,

$$\sup_{x,y\in C_{j,k}\cap \mathrm{supp}(\Pi)} \|x - y\|_2 \leq \theta_2 2^{-j}.$$

(A3) Let $\lambda_1^{j,k}, \ldots, \lambda_D^{j,k}$ denote the eigenvalues of the covariance matrix $\Sigma_{j,k}$ (defined in 2.1) arranged in the non-increasing order. Then there exist positive constants $\sigma = \sigma(\Pi)$, $\theta_3 = \theta_3(\Pi)$, $\theta_4 = \theta_4(\Pi)$, and some $\alpha \in (0, 1]$ such that for all $k = 1 \ldots N(j)$, $\lambda_d^{j,k} \geq \theta_3 2^{-2j}/d$ and

$$\sum_{l=d+1}^{D} \lambda_l^{j,k} \leq \theta_4(\sigma^2 + 2^{-2(1+\alpha)j}) \leq \frac{1}{2}\lambda_d^{j,k}.$$

Theorem 3.1 ([20]). *Suppose (A1), (A2), and (A3), and let $X_1, \ldots, X_n$ be an i.i.d. sample from Π, and set $\bar{d} := 4d^2\theta_2^4/\theta_3^2$. Then for any $j_{\min} \leq j \leq j_{\max}$ and any $t \geq 1$ such that $t + \log(\bar{d} \vee 8) \leq \frac{1}{2}\theta_1 n 2^{-jd}$,*

$$\begin{aligned}
\mathbb{E}\|x - \widehat{P}_j(x)\|_2^2 \leq\ & 2\theta_4\left(\sigma^2 + 2^{-2j(1+\alpha)}\right)\\
&+ c_1 2^{-2j}\frac{(t + \log(\bar{d} \vee 8))d^2}{n2^{-jd}}
\end{aligned}$$

with probability $\geq 1 - \frac{2^{jd+1}}{\theta_1}\left(e^{-t} + e^{-\frac{\theta_1}{16}n2^{-jd}}\right)$, where $c_1 = 2\left(12\sqrt{2}\frac{\theta_2^3}{\theta_3\sqrt{\theta_1}} + 4\sqrt{2}\frac{\theta_2}{d\sqrt{\theta_1}}\right)^2$.

3.2 Special case: noisy manifolds

We now show that there is a large class of interesting distributions Π and a data-dependent partitioning algorithm such that (A1), (A2), and (A3) are satisfied with high probability given a sufficiently dense sample.

Noisy manifold models: Let $\mathcal{M}$ be a smooth, compact d-dimensional submanifold of $\mathbb{R}^D$. We recall the definition of the *reach* [16], an important global characteristic of $\mathcal{M}$. Let $D(\mathcal{M}) = \{y \in \mathbb{R}^D : \exists! x \in \mathcal{M} \text{ s.t. } \|x - y\|_2 = \inf_{z\in\mathcal{M}} \|z - y\|_2\}$, $\mathcal{M}_r = \{y \in \mathbb{R}^D : \inf_{x\in\mathcal{M}} \|x - y\|_2 < r\}$. Then $\mathrm{reach}(\mathcal{M}) := \sup\{r \geq 0 : \mathcal{M}_r \subseteq D(\mathcal{M})\}$. Assume that $0 < \sigma < \tau$. We shall say that the distribution Π satisfies the (τ, σ)-model assumption if there exist a compact, smooth submanifold $\mathcal{M} \hookrightarrow \mathbb{R}^D$ without boundary and with reach τ such that $\mathrm{supp}(\Pi) = \mathcal{M}_\sigma$, Π is absolutely continuous with respect to $U_{\mathcal{M}_\sigma}$ - the uniform distribution on $\mathcal{M}_\sigma$, and the Radon-Nikodym derivative $\frac{d\Pi}{dU_{\mathcal{M}_\sigma}}$ satisfies

$$0 < \phi_1 \leq \frac{d\Pi}{dU_{\mathcal{M}_\sigma}} \leq \phi_2 < \infty$$

$\mathcal{M}_\sigma$ - almost surely.

Data-dependent partitions: Our partitioning scheme is based on the *cover trees* [4] data structure. We will construct a cover tree for the collection $S = \{X_1, \ldots, X_n\}$ of i.i.d. samples from the distribution Π with respect to the Euclidean distance $d(x, y) := \|x - y\|_2$. Assume that $T_j := T_j(X_1, \ldots, X_n) = \{a_{j,k}\}_{k=1}^{N(j)}$ is the set of nodes on level j of the cover tree T for S. Define the indexing map $k(x) := \operatorname{argmin}_{1\leq k\leq N(j)} \|x - a_{j,k}\|_2$ (the ties are broken by choosing the smallest value of k), and partition $\mathbb{R}^D$ into the Voronoi regions $C_{j,k} = \{x \in \mathbb{R}^D : k_j(x) = k\}$. Let $j(n, t)$ be the maximal $j \in \mathbb{Z}_+$ such that $n \geq C(\Pi)2^{jd}(jd + t)$.

Theorem 3.2. *Assume that Π satisfies the (τ, σ)-model assumption with $\sigma \leq c_2(d)\tau$, and let $X_1, \ldots, X_n$ be an i.i.d. sample from Π. Suppose that $C_{j,k}$ are defined as above. Then there exist $c_3, c_4 > 0$ depending only on d such that with probability $\geq 1 - e^{-t}$, $C_{j,k}$ satisfy (A1), (A2), and (A3) for $j_{\min} \geq \log_2(c_3/\tau)$, $j_{\max} \leq \log_2(c_4/\sigma) \wedge j(n, t)$ with $\alpha = 1$, $\sigma(\Pi) := \sigma$ and θ_1, θ_2, θ_3, θ_4 which are independent of j and D.*

4 Discussion

Theorem 3.2 demonstrates that distributions supported on noisy manifolds interact in an inherently low-dimensional manner with partitions obtained via the cover trees algorithm. This is true as long as the resolution level of the partition remains above the noise level.

Theorem 3.1 supplies us with an error bound for the empirical GMRA provided that the distribution Π and the partition $\{C_{j,k}\}$ interact in an inherently low-dimensional manner. Pointwise bounds may also be obtained. This gives a rate in the noiseless case comparable to the minimax rate in [8], but with respect to the Hausdorff distance. A detailed discussion is left to [20].

For practical purposes, Theorem 3.2 may be combined with Theorem 3.1 as follows: given an i.i.d. sample $\{X_1, \ldots, X_{2n}\}$ from Π, use the first n points $\{X_1, \ldots, X_n\}$ to construct the partition $C_{j,k}$ using the cover tree algorithm, and then the remaining $\{X_{n+1}, \ldots, X_{2n}\}$ are used to obtain $\widehat{P}_j$ (see 2.2). This makes our method entirely data-dependent.

References

[1] M. Aharon, M. Elad, and A. Bruckstein. K-SVD: An algorithm for designing overcomplete dictionaries for sparse representation. *Signal Processing, IEEE Transactions on [see also Acoustics, Speech, and Signal Processing, IEEE Transactions on]*, 54(11):4311–4322, 2006.

[2] W. K. Allard, G. Chen, and M. Maggioni. Multi-scale geometric methods for data sets II: Geometric multi-resolution analysis. *Applied and Computational Harmonic Analysis*, 32(3):435–462, 2012. (submitted:5/2011).

[3] M. Belkin and P. Niyogi. Using manifold structure for partially labelled classification. *Advances in NIPS*, 15, 2003.

[4] A. Beygelzimer, S. Kakade, and J. Langford. Cover trees for nearest neighbor. In *ICML*, pages 97–104, 2006.

[5] G. Chen and M. Maggioni. Multiscale geometric and spectral analysis of plane arrangements. In *Proc. CVPR*, 2011.

[6] G. Chen and M.Maggioni. Multiscale geometric wavelets for the analysis of point clouds. *To appear in Proc. CISS 2010*, 2010.

[7] S. S. Chen, D. L. Donoho, and M. A. Saunders. Atomic decomposition by basis pursuit. *SIAM Journal on Scientific Computing*, 20(1):33–61, 1998.

[8] I. V. Christopher R. Genovese, Marco Perone-Pacifico and L. Wasserman. Manifold estimation and singular deconvolution under hausdorff loss. *Ann. Statist.*, 40(2):941–963, 2012.

[9] R. Coifman and M. Maggioni. Diffusion wavelets. *Appl. Comp. Harm. Anal.*, 21(1):53–94, July 2006. (Tech. Rep. YALE/DCS/TR-1303, Yale Univ., Sep. 2004).

[10] R. R. Coifman, S. Lafon, A. B. Lee, M. Maggioni, B. Nadler, F. Warner, and S. W. Zucker. Geometric diffusions as a tool for harmonic analysis and structure definition of data: Diffusion maps. *PNAS*, 102(21):7426–7431, 2005.

[11] R. R. Coifman, S. Lafon, A. B. Lee, M. Maggioni, B. Nadler, F. Warner, and S. W. Zucker. Geometric diffusions as a tool for harmonic analysis and structure definition of data: Multiscale methods. *PNAS*, 102(21):7432–7438, 2005.

[12] D. Donoho. Compressed sensing. *IEEE Tran. on Information Theory*, 52(4):1289–1306, April 2006.

[13] D. L. Donoho and C. Grimes. When does isomap recover natural parameterization of families of articulated images? Technical Report 2002-27, Department of Statistics, Stanford University, August 2002.

[14] D. L. Donoho and C. Grimes. Hessian eigenmaps: new locally linear embedding techniques for high-dimensional data. *Proc. Nat. Acad. Sciences*, pages 5591–5596, March 2003.

[15] E. Elhamifar and R. Vidal. Sparse subspace clustering. In *CVPR*, pages 2790–2797, June 2009.

[16] H. Federer. Curvature measures. *Transactions of the American Mathematical Society*, 93(3):418–491, 1959.

[17] H. Hotelling. Analysis of a complex of statistical variables into principal components. *Journal of Educational Psychology*, 24(4):17–441,498–520, 1933.

[18] P. Jones, M. Maggioni, and R. Schul. Manifold parametrizations by eigenfunctions of the Laplacian and heat kernels. *Proc. Nat. Acad. Sci.*, 105(6):1803–1808, Feb. 2008.

[19] P. Jones, M. Maggioni, and R. Schul. Universal local manifold parametrizations via heat kernels and eigenfunctions of the Laplacian. *Ann. Acad. Scient. Fen.*, 35:1–44, January 2010. http://arxiv.org/abs/0709.1975.

[20] M. Maggioni, S. Minsker, and N. Strawn. Dictionary learning and non-asymptotic bounds for Geometric Multi-Resolution Analysis. *Preprint*, 2014.

[21] K. Pearson. On lines and planes of closest fit to systems of points in space. *Philosophical Magazine*, 2(11):559–572, 1901.

[22] S. Roweis and L. Saul. Nonlinear dimensionality reduction by locally linear embedding. *Science*, 290:2323–2326, 2000.

[23] Y. Sugaya and K. Kanatani. Multi-stage unsupervised learning for multi-body motion segmentation. *IEICE Transactions on Information and Systems*, E87-D(7):1935–1942, 2004.

[24] J. B. Tenenbaum, V. D. Silva, and J. C. Langford. A global geometric framework for nonlinear dimensionality reduction. *Science*, 290(5500):2319–2323, 2000.

[25] R. Vidal. Subspace clustering. *IEEE Signal Processing Magazine*, 28(2):52–68, 2011.

[26] R. Vidal, Y. Ma, and S. Sastry. Generalized principal component analysis (GPCA). *PAMI*, 27(12), 2005.

[27] T. Zhang, A. Szlam, Y. Wang, and G. Lerman. Randomized hybrid linear modeling by local best-fit flats. In *CVPR*, San Francisco, CA, June 2010.

[28] Z. Zhang and H. Zha. Principal manifolds and nonlinear dimension reduction via local tangent space alignment. *SIAM Journal of Scientific Computing*, 26:313–338, 2002.

An application of the Chambolle-Pock algorithm to 3D + time tomography

Cyril Mory[1,2] and Laurent Jacques[3].

[1]iMagX project, ICTEAM Institute, Université catholique de Louvain, Louvain-la-Neuve, Belgium.

[2]Université de Lyon, CREATIS ; CNRS UMR5220 ; Inserm U1044 ; INSA-Lyon ; Université Lyon 1 ; Centre Léon Bérard, France.

[3]ISP GROUP, ICTEAM/ELEN, Université catholique de Louvain, Louvain-la-Neuve, Belgium.

***Abstract*— The 4D RecOnstructiOn using Spatial and TEmporal Regularization method is a recent 4D cone beam computed tomography algorithm. 4D ROOSTER has not been rigorously proved to converge. This paper aims to reformulate it using the Chambolle & Pock primal-dual optimization scheme. The convergence of this reformulated 4D ROOSTER is therefore guaranteed.**

1 Introduction

Four dimensional cone beam computed tomography (4D CBCT) of the free breathing thorax is important for image-guided radiation therapy (IGRT). Mainly two families of methods have been proposed to handle the problem: respiratory motion compensation, in which the motion of organs during breathing is estimated and used to deform the volume during reconstruction, and respiration-correlated reconstruction, in which several volumes are computed, each one from a subset of the projections that has been acquired during the correct respiratory phase. Advances in compressed sensing have been used to improve respiration-correlated reconstruction by enforcing spatial regularity constraints [1, 2, 3, 4], but only a few recent methods exploit the strong correlation between successive respiratory phases [5, 6, 7, 8]. 4D ROOSTER [8], which was recently proposed for 4D cardiac CBCT, is one such method. However, it came without a rigorous proof of convergence. In this paper, we propose a reformulation of 4D ROOSTER using the Chambolle & Pock primal-dual optimization scheme [9]. The convergence of this reformulated 4D ROOSTER is therefore guaranteed.

2 The original 4D ROOSTER method

This algorithm assumes that a rough segmentation of the patient's heart is available, and that movement is expected to occur only inside this segmented region. The method consists in iteratively enforcing five different constraints in an alternating manner. It starts by minimizing a quadratic data-attachment term $\sum_{\alpha} \|R_\alpha S_\alpha x - p_\alpha\|_2^2$, with α the projection angle, x a 4D sequence of volumes, R_α the forward projection operator at angle α, S_α a linear interpolator, and p_α the measured projection at angle α. This data-attachment term is minimized by conjugate gradient. Then the following regularization steps are applied sequentially: positivity enforcement, averaging along time outside the heart region, spatial total-variation denoising, and temporal total-variation denoising. This constitutes one iteration of the main loop, the output of which is fed back to the conjugate gradient minimizer for the next iteration. This algorithm offers no convergence guarantees. In [10, 11], using the theory of non-expansive mappings, it is only proved that if the main iteration has at least one fixed point, the algorithm converges to one of them. However, we show in this paper that each step of this method can be interpreted as a proximal operator [8], thus little effort is required to make it fit into the Chambolle & Pock framework. We remind that for $f : \mathbb{R}^N \to \mathbb{R}$ a closed convex function, the proximal operator of f is defined as $\text{prox}_f(v) = \underset{x \in \mathbb{R}^N}{\arg\min}\ \frac{1}{2}\|x - v\|_2^2 + f(x)$.

3 Chambolle & Pock 4D ROOSTER

Setting $R = \begin{pmatrix} R_{\alpha_1} S_{\alpha_1} \\ \vdots \\ R_{\alpha_m} S_{\alpha_m} \end{pmatrix}$ and $p = \begin{pmatrix} p_{\alpha_1} \\ \vdots \\ p_{\alpha_m} \end{pmatrix}$, with m the number of projections, the data-attachment term becomes a single L^2 norm $\sum_{\alpha} \|R_\alpha S_\alpha x - p_\alpha\|_2^2 = \|Rx - p\|_2^2$. The 4D ROOSTER optimization problem can then be expressed as the search for

$$\begin{aligned} \underset{x \in \mathbb{R}^N}{\arg\min}\ & \frac{1}{2}\|Rx - p\|_2^2 + \lambda_2 \|x\|_{TV_{space}} \\ & + \lambda_3 \|x\|_{TV_{time}} + i_{\mathbb{R}^{N+}}(x) + i_{ROI}(x) \end{aligned} \tag{1}$$

with x the 3D+t sequence of volumes, with in total N voxels, p the set of measured projections, with in total P pixels, $R : \mathbb{R}^N \to \mathbb{R}^P$ the forward projection operator, $\|.\|_{TV_{space}}$ the spatial total-variation norm, $\|.\|_{TV_{time}}$ the temporal total-variation norm, $\mathbb{R}^{N+}$ the set of sequences of volumes in which all voxels have non-negative values, ROI the set of sequences of volumes in which all voxels outside the heart have equal values, $i_{\mathbb{R}^{N+}}(x)$ and $i_{ROI}(x)$ their respective convex indicator functions, and $\lambda_2, \lambda_3 > 0$ two parameters weighing the relative importance of the terms. We adopt the same formalism as in [12], and therefore use the following notations:

$$\begin{aligned} & F_1 : \mathbb{R}^P \to \mathbb{R}, t \to \frac{1}{2}\|t - p\|_2^2 \\ & F_2 : \mathbb{R}^N \to \mathbb{R}, t \to \lambda_2 \|t\|_{TV_{space}} \\ & F_3 : \mathbb{R}^N \to \mathbb{R}, t \to \lambda_3 \|t\|_{TV_{time}} \\ & H : \mathbb{R}^N \to \mathbb{R} \cup \{+\infty\}, t \to i_{\mathbb{R}^{N+}}(t) + i_{ROI}(t) \end{aligned}$$

Because the cost function has more than two terms, we have to reformulate the problem into a search in $\mathbb{R}^{3N}$ by defining $x' = \begin{pmatrix} x_1 \\ x_2 \\ x_3 \end{pmatrix}$, with x_1, x_2 and $x_3 \in \mathbb{R}^N$, and

$\prod_{1,i} = \{x' \in \mathbb{R}^{3N} \mid x_1 = x_i\}$ for $i = 2, 3$. The 4D ROOSTER optimization problem becomes

$$\begin{aligned}\arg\min_{x\in\mathbb{R}^{3N}} & F_1(Rx_1) + F_2(x_2) + F_3(x_3) \\ & + i_{\mathbb{R}^{N+}}(x_1) + i_{ROI}(x_1) + i_{\prod_{1,2}}(x') + i_{\prod_{1,3}}(x')\end{aligned} \tag{2}$$

We define the total dimension $W = P + 2N$, the convex functions F and G and the linear operator K:

$$\begin{aligned} F &: \mathbb{R}^W \to \mathbb{R}, s = \begin{pmatrix} s_1 \\ s_2 \\ s_3 \end{pmatrix} \to \sum_{j=1..3} F_j(s_j) \\ &\text{with } s_1 \in \mathbb{R}^P, s_2 \in \mathbb{R}^N, s_3 \in \mathbb{R}^N \\ K &: \mathbb{R}^{3N} \to \mathbb{R}^W, K = \begin{pmatrix} R & 0 & 0 \\ 0 & I_N & 0 \\ 0 & 0 & I_N \end{pmatrix} \\ G &: \mathbb{R}^{3N} \to \mathbb{R}, x' \to i_{\prod_{1,2}}(x') + i_{\prod_{1,3}}(x') + H(x_1) \end{aligned}$$

With these new notations, we can write the cost function as a sum of two functions $\arg\min_{x\in\mathbb{R}^{3N}} F(Kx') + G(x')$ on which the Chambolle-Pock algorithm applies. In compact form, the Chambolle-Pock algorithm is as follows: let $v'^{(0)} \in \mathbb{R}^W, x'^{(0)} \in \mathbb{R}^{3N}, y'^{(0)} \in \mathbb{R}^{3N}$, the update step is

$$\begin{cases} v'^{(k+1)} = \text{prox}_{\gamma F^*}\left(v'^{(k)} + \gamma K y'^{(k)}\right) \\ x'^{(k+1)} = \text{prox}_{\mu G}\left(x'^{(k)} - \mu K^* v'^{(k+1)}\right) \\ y'^{(k+1)} = 2x'^{(k+1)} - x'^{(k)} \end{cases}$$

If we expand this for our problem and simplify it since $x_1 = x_2 = x_3$ and $y_1 = y_2 = y_3$, we obtain:

$$\begin{cases} v_1^{(k+1)} = \text{prox}_{\gamma F_1^*}\left(v_1^{(k)} + \gamma R y_1^{(k)}\right) \to \text{Data fidelity} \\ v_2^{(k+1)} = \text{prox}_{\gamma F_2^*}\left(v_2^{(k)} + \gamma y_2^{(k)}\right) \to \text{TV in space} \\ v_3^{(k+1)} = \text{prox}_{\gamma F_3^*}\left(v_3^{(k)} + \gamma y_3^{(k)}\right) \to \text{TV in time} \\ x^{(k+1)} = \text{prox}_{\frac{\mu}{3}H}\left(3x_1^{(k)} - \mu\left(R^* v_1^{(k+1)}\right.\right. \\ \left.\left. + v_2^{(k+1)} + v_3^{(k+1)}\right)\right) \to \text{Positivity and ROI} \\ y^{(k+1)} = 2x^{(k+1)} - x^{(k)} \to \text{Update step} \end{cases}$$

Four new operators appear in this formulation. $R^* : \mathbb{R}^P \to \mathbb{R}^N$ is the back projection operator. To express the other ones, we use the relation between the proximal operator of a function and the proximal operator of its adjoint:

$$\text{prox}_{\gamma F^*}(x) = x - \gamma \,\text{prox}_{\frac{1}{\gamma}F}\left(\frac{1}{\gamma}x\right)$$

Therefore

$$\text{prox}_{\gamma F_1^*}(x) = x - \gamma\left(\frac{x+p}{\gamma+1}\right) = \frac{x-\gamma p}{\gamma+1}$$

The proximal operator for the spatial and temporal TV norms can be computed iteratively using the Chambolle algorithm [13], simply by using either the spatial or the temporal gradient and divergence operators

$$\begin{aligned} &\text{prox}_{\frac{1}{\gamma}TV}(x) = x - \tfrac{1}{\gamma}\,\text{div}\,\overline{p}, \text{where } \overline{p} = \lim_{n\to+\infty} p^{(n)} \\ &p^{(0)} = 0 \text{ and for every voxel } i, \\ &p_i^{(n+1)} = \frac{p_i^n + h\left(\nabla\left(\text{div}\, p^{(n)} - \gamma x\right)\right)_i}{1 + h\left(\nabla\left(\text{div}\, p^{(n)} - \gamma x\right)\right)_i}, \text{with } h > 0 \end{aligned}$$

And the proximal operator of γF_2^*(and similarly γF_3^*) can be obtained from the same relationship as previously

$$\begin{aligned} &\text{prox}_{\gamma F_2^*}\left(v_2^{(k)} + \gamma y_2^{(k)}\right) \\ &\quad = v_2^{(k)} + \gamma y_2^{(k)} - \gamma\,\text{prox}_{\frac{1}{\gamma}F_2}\left(\frac{1}{\gamma}\left(v_2^{(k)} + \gamma y_2^{(k)}\right)\right) \\ &\quad = v_2^{(k)} + \gamma\left(y_2^{(k)} - \text{prox}_{\frac{1}{\gamma}F_2}\left(\frac{1}{\gamma}v_2^{(k)} + y_2^{(k)}\right)\right) \end{aligned}$$

With this new formulation, the modified 4D ROOSTER is guaranteed to converge. The theory of proximal algorithms also helps find the suitable parameters based on the norm of R, while they can only be determined empirically in the original method.

4 Perspectives

4D ROOSTER can be applied to respiration-correlated 4D CBCT in order to reduce the level of streak artifacts in the reconstructions. Preliminary results are provided in Figure 1 and 2. They have been obtained using the original 4D ROOSTER method, and are compared with a standard ECG-gated FDK reconstruction. The modified 4D ROOSTER method presented in this paper will be implemented in the near future using the RTK framework [14].

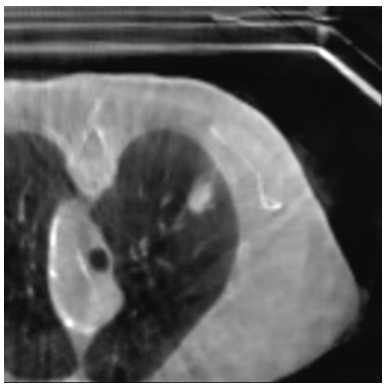

Figure 1: Respiratory-gated reconstruction of a lung tumor using 4D ROOSTER

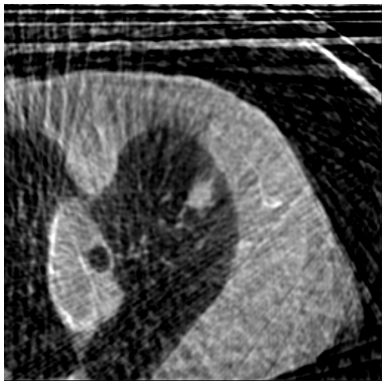

Figure 2: Respiratory-gated reconstruction of a lung tumor using the Feldkamp, Davis and Kress method

References

[1] F. Bergner, T. Berkus, M. Oelhafen, P. Kunz, T. Pan, R. Grimmer, L. Ritschl, and M. Kachelrieß, "An investigation of 4D cone-beam CT algorithms for slowly rotating scanners," *Medical Physics*, vol. 37, no. 9, p. 5044, 2010.

[2] Z. Qi and G.-H. Chen, "Extraction of tumor motion trajectories using PICCS-4DCBCT: a validation study," *Medical physics*, vol. 38, no. 10, pp. 5530–5538, Oct. 2011, PMID: 21992371 PMCID: PMC3195374.

[3] S. Leng, J. Tang, J. Zambelli, B. Nett, R. Tolakanahalli, and G.-H. Chen, "High temporal resolution and streak-free four-dimensional cone-beam computed tomography," *Physics in medicine and biology*, vol. 53, no. 20, pp. 5653–5673, Oct. 2008, PMID: 18812650.

[4] E. Y. Sidky and X. Pan, "Image reconstruction in circular cone-beam computed tomography by constrained, total-variation minimization," *Physics in Medicine and Biology*, vol. 53, no. 17, pp. 4777–4807, Sep. 2008.

[5] H. Gao, R. Li, Y. Lin, and L. Xing, "4D cone beam CT via spatiotemporal tensor framelet," *Medical Physics*, vol. 39, no. 11, pp. 6943–6946, Nov. 2012, PMID: 23127087 PMCID: PMC3494730.

[6] H. Wu, A. Maier, R. Fahrig, and J. Hornegger, "Spatial-temporal total variation regularization (STTVR) for 4D-CT reconstruction," in *Proceedings of SPIE Medical Imaging 2012*, N. J. Pelc, R. M. Nishikawa, and B. R. Whiting, Eds., San Diego, CA, USA, Feb. 2012, p. 83133J.

[7] Z. Tian, X. Jia, B. Dong, Y. Lou, and S. B. Jiang, "Low-dose 4DCT reconstruction via temporal nonlocal means," *Medical physics*, vol. 38, no. 3, pp. 1359–1365, Mar. 2011, PMID: 21520846.

[8] C. Mory, "Cardiac c-arm computed tomography," Ph.D. dissertation, Université Lyon 1, Feb. 2014.

[9] A. Chambolle and T. Pock, "A first-order primal-dual algorithm for convex problems with applications to imaging," *Journal of Mathematical Imaging and Vision*, vol. 40, no. 1, pp. 120–145, May 2011.

[10] J. J. Moreau, "Fonctions convexes duales et points proximaux dans un espace hilbertien," *Comptes Rendus de l'Académie des Sciences (Paris)*, vol. 255, pp. 2897–2899, 1962.

[11] P. L. Combettes and J.-C. Pesquet, "Proximal splitting methods in signal processing," in *Fixed-Point Algorithms for Inverse Problems in Science and Engineering*, ser. Springer Optimization and Its Applications, H. H. Bauschke, R. S. Burachik, P. L. Combettes, V. Elser, D. R. Luke, and H. Wolkowicz, Eds. Springer New York, Jan. 2011, pp. 185–212.

[12] A. Gonzalez, L. Jacques, C. De Vleeschouwer, and P. Antoine, "Compressive optical deflectometric tomography: A constrained total-variation minimization approach," *Journal of Inverse Problems and Imaging*, In press.

[13] A. Chambolle, "An algorithm for total variation minimization and applications," *J. Math. Imaging Vis.*, vol. 20, no. 1-2, p. 89–97, Jan. 2004.

[14] S. Rit, M. Vila Oliva, S. Broumische, R. Labarbe, D. Sarrut, and G. C. Sharp, "The reconstruction toolkit (RTK), an open-source cone-beam CT reconstruction toolkit based on the insight toolkit (ITK)," in *Proceedings of the International Conference on the Use of Computers in Radiation Therapy (ICCR)*, 2013.

Super-resolution method using sparse regularization for point spread function recovery

F. Ngole[1] and J.-L. Starck[2].

[1,2]Laboratoire AIM, UMR CEA-CNRS-Paris 7, Irfu, Service d'Astrophysique, CEA Saclay, F-91191 GIF-SUR-YVETTE Cedex, France.

***Abstract*— In large scale spatial surveys such as the forthcoming ESA Euclid mission, images may be undersampled due to the optical sensors sizes. Therefore, one may consider using a super-resolution (SR) method to recover aliased frequencies, prior to further analysis. This is particularly relevant for point source images which provide direct measurements of the instrument point spread function (PSF). Super-resolution is a classical problem in image processing literature. It can be formulated as an inverse problem and sparsity constraint may be applied. We show that using an appropriate wavelet dictionary, this provides significant improvements over existing super-resolution methods, especially on the PSF.**

1 Introduction

The weak lensing is one of the most promising tools to probe the dark matter distribution in the universe. The idea is to infer, from billion galaxies images, the shapes distortions due to dark matter gravitational lensing and then constraint the dark matter mass distribution. The Euclid mission will provide the data for such a purpose [1]. Nevertheless, galaxies images will be distorted due to the telescope PSF. It's critical to take accurately this PSF into account. In practice, isolated stars provide measurements of the PSF at different locations in the field of view. But these stars images will be aliased in Euclid, given the CCD sensor sizes. Nevertheless, several low resolution (LR) versions of the same PSF will be available thanks to the multiple exposures planned during the survey. This naturally leads to consider a SR approach to restore an accurate PSF.

SR techniques generally adopt the following observation model. We suppose that the desired high resolution (HR) image of size $d_1p_1 \times d_2p_2$ is written in lexicographic order in a vector $\mathbf{x} = (x_1, ..., x_q)^T$ where $q = d_1p_1d_2p_2$ and d_1 and d_2 are respectively the line and column downsampling factors. Let $\mathbf{y_k} = (y_{k1}, ..., y_{kp})^T$ denote the k^{th} LR observation written in lexicographical order, with $p = p_1p_2$ and $k = 1...n$. So we have:

$$\mathbf{y_k} = \mathbf{DB_kM_kx} + \mathbf{n_k},\ k = 1...n \tag{1}$$

where $\mathbf{M_k}$ is a warp matrix, $\mathbf{B_k}$ is a blur matrix, $\mathbf{D_k}$ is a decimation matrix and $\mathbf{n_k}$ is a noise column vector. The matrices $\mathbf{M_k}$ traduce the motions of the observations relatively to each others and generally need to be estimated. The matrices $\mathbf{B_k}$ account for different blurs (the atmosphere blur, which is particularly considerable for ground based telescopes, the system optics blur, the imaging system shaking, pixel response etc.). And $\mathbf{D}$ is simply a downsampling operator which is assumed to be the same for all the observations.

In our case, we are interested in estimating the telescope PSF, in other terms the telescope contribution to the blur. We will consider that the telescope is steady so that the blur is stable in time. The Euclid telescope will be out of the atmosphere so that under our hypothesis, the blur only consists of the telescope PSF. Thus, the index k may be dropped for the blur operator. Moreover, we will consider that the PSF varies slowly in the field so that for small motions between observations, the warp and blur matrix may be commuted. So we can rewrite the model 1 as

$$\mathbf{y_k} = \mathbf{W_k}\tilde{\mathbf{x}} + \mathbf{n_k},\ k = 1...n \tag{2}$$

where $\tilde{\mathbf{x}} = \mathbf{Bx}$ and $\mathbf{W_k} = \mathbf{DM_k}$ and we will be interested in recovering $\tilde{\mathbf{x}}$ with $\mathbf{x}$ being assumed to be a point source.

In the next section, we will describe some SR techniques dedicated to astronomical images, but one may refer to [2] for more details on SR various framework in general image processing literature.

2 SR techniques in astronomy

The most simple super-resolution method is certainly the shift-and-add method. It performs in three steps : first the images are upsampled to the target resolution. Then they are shifted on a common grid and averaged. It has been used in astronomy for a long time, particularly for ground based telescopes. Despite its simplicity, it has been shown that this method provides an optimal solution in the sense of the maximum likelihood (ML) with additive white gaussian noise (WGN) and when only pure translation motions are considered. The interpolation operator should be $\mathbf{D}^T$ (from Eq. 1) which comes down to a simple zero-padding interpolation, and the averaging should be performed with proper weighting[3].

Now let consider the SR method implemented in PSFEXTRACTOR, an open source software more and more used in astronomy community for PSF modeling. This software and the associated documentation may be found on the website `http://www.astromatic.net/`.The desired HR image is now a matrix $\mathbf{X}$ of size $dp \times dp$ and the k^{th} observation $\mathbf{Y_k}$ is a matrix of size $p \times p$ with $k = 1...n$. The coordinates of the centroid of the k^{th} observation are noted (i_k, j_k). Assuming that the images are bandlimited, the samples of the LR images can be interpolated from the desired HR image, thanks to the well-known Shannon sampling theorem. In theory, this interpolation should involve a 2D sine cardinal kernel which will be approximated with a finite support function. Let $h(.,.)$ denote such a function. The estimate of the sample (i, j) of the k^{th} observation is given by:

$$\hat{y}_{k,ij} = \sum_l \sum_m h\left[l - d(i - i_k), m - d(j - j_k)\right] x_{ij}. \tag{3}$$

Then we can define the cost function

$$J_1(\mathbf{X}) = \sum_{k=1}^{n} \sum_{i=1}^{p} \sum_{j=1}^{p} \frac{(y_{k,ij} - f_k\hat{y}_{k,ij})^2}{\sigma_k^2} \tag{4}$$

where f_k accounts for possible luminosity differences and needs to be estimated as well as σ_k^2 which is related to the local background variance and the pixel values uncertainty.

PSFEXTRACTOR uses a Lanczos4 interpolant for h. Finally, PSFEXTRACTOR minimizes the cost function defined as

$$J_2(\boldsymbol{\Delta}) = J_1(\boldsymbol{\Delta} + \mathbf{X}^{(0)}) + \lambda \|\boldsymbol{\Delta}\|_{l_2}^2 \tag{5}$$

where $\boldsymbol{\Delta} = \mathbf{X} - \mathbf{X}^{(0)}$ and $\mathbf{X}^{(0)}$ is an median image computed from the LR observations. This second term is meant for regularization purposes in ill-conditioned or ill-posed cases ($n < d^2$).

3 Sparse regularization

Now we introduce the SPRITE algorithm(SParse Recovery of InsTrumental rEsponse) which as stated before basically consists in adding a sparse penalty in the recovery process. As previously, our optimization variable will be the difference between the desired HR image and a first guess. But instead of using a quadratic regularization term, we impose the analysis coefficients of this difference to be sparse in a chosen wavelet dictionary. Let note $\boldsymbol{\Phi}$ such a dictionary so that the wavelet coefficients are given by $\boldsymbol{\alpha_x} = \boldsymbol{\Phi}\mathbf{x}$, then such a prior can be integrated by solving

$$\min_{\boldsymbol{\Delta}} J_1(\boldsymbol{\Delta} + \mathbf{X}^{(0)}) + \gamma \|\boldsymbol{\Phi}\boldsymbol{\Delta}\|_{l_1}, \tag{6}$$

with $\boldsymbol{\Delta}$ defined as in 5 and $\mathbf{X}^{(0)}$ calculated using a shift-and-add. A similar l_1 penalty has already been applied for SR. An example may be found in [4], where the cost function is minimized using variants of the alternating direction method of multipliers (ADMM). Besides advantages of such approaches over quadratic regularizers have been proved in many related problems in image and signal processing. It has been shown in [5] that under certain conditions on $\boldsymbol{\Phi}$ the proximity operator of $\|.\|_{l_1} \circ \boldsymbol{\Phi}$ is well defined and can be estimated iteratively so that a forward-backward algorithm can be used to solve 6 ([6]). Nevertheless, we opted for a simple iterative shrinkage scheme. Let $\hat{\boldsymbol{\Phi}}$ denote a reconstruction operator associated with $\boldsymbol{\Phi}$. The optimization variable is updated as follows:

$$\boldsymbol{\Delta_{n+1}} = \hat{\boldsymbol{\Phi}}\, \mathrm{SoftThresh}_{\gamma} \left[\boldsymbol{\Phi}(\boldsymbol{\Delta_n} - \mu(\nabla J_1(\boldsymbol{\Delta_n} + \mathbf{X}^{(0)}) \right]. \tag{7}$$

Proceeding this way amounts to alternately minimizing the $J_1(\boldsymbol{\Delta}+\mathbf{X}^{(0)})$ and $\gamma\|\boldsymbol{\alpha_\Delta}\|_{l_1}$ with $\boldsymbol{\alpha_\Delta} = \boldsymbol{\Phi}\boldsymbol{\Delta}$. This approach is inspired from the MCA decomposition algorithm in [7]. In our experiments, this gave qualitatively equivalent results as when we used the rigorously estimated analysis proximal operator, with a radical gain in complexity. Furthermore, we used the scheme introduced in [8] for FISTA to accelerate the convergence. In practice, the parameter γ, which is used as a threshold, is set at 3σ where σ is the estimated noise level at each considered wavelet scale. We used an undecimated wavelet transform with a 7/9 biorthogonal filter bank which is quite convenient since the data are nearly isotropic. Besides, we observe that reweighting the sparse penalty term as suggested in [9] tends in general to improve the final solution. Thus, the algorithm includes one "reweighted" iteration in which for each wavelet coefficient $\boldsymbol{\alpha_{\Delta i}}$, the threshold γ is now replaced with $\frac{\gamma}{1+|\boldsymbol{\alpha}_{\boldsymbol{\Delta i}}^{(0)}|/3\sigma}$, where $\boldsymbol{\alpha}_{\boldsymbol{\Delta i}}^{(0)}$ is the wavelet coefficient previously calculated at the position i after thresholding. Thus in the reweighted iteration the thresholds will be lower on the stronger coefficients.

4 Numerical experiments

For Euclid mission, the actual sampling frequency is about 0.688 times the Nyquist frequency([10]). Therefore we will target an upsampling factor of 2, which gives a sufficient bandpass to recover the high frequencies. Our data set consists of simulated optical point spread functions generated with the software Zemax, using Euclid telescope parameters. Our data set contains simulated PSF measurements on a regular 18×18 grid on the field of view (324 PSF). The original PSF are downsampled to almost Nyquist frequency. For each PSF, 4 randomly shifted "copies" are generated and downsampled with a factor 2 in lines and columns to achieve Euclid resolution . Different levels of WGN are added. We define the signal level as its empirical variance calculated in a 50×50 patch centered on the PSF main lobe. For each algorithm, we used these 4 images to reconstruct the original PSF. For quality assessment, we used the average error between the reference PSF ellipticity parameters and the super-resolved PSF ellipticity using each algorithm and these errors standard deviations. These ellipticity parameters definitions can be found in [10] and they are critical in weak lensing measurements.

As it can be seen on the figures below, smaller errors and errors dispersion are achieved with SPRITE algorithm, especially at low SNR.

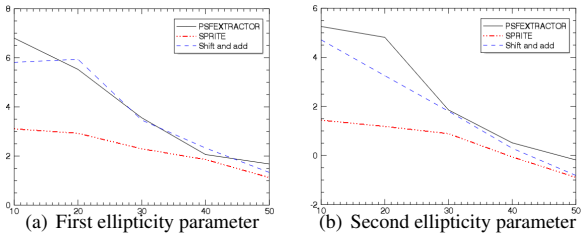

(a) First ellipticity parameter (b) Second ellipticity parameter

Figure 1: Average absolute values (in log) of relative errors on ellipticity parameters versus the SNR

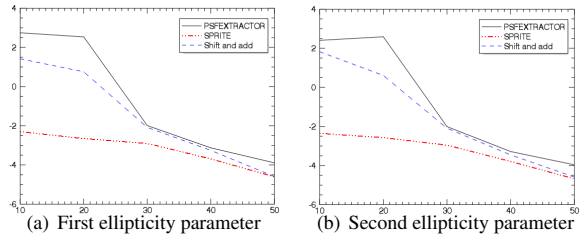

(a) First ellipticity parameter (b) Second ellipticity parameter

Figure 2: Average standard deviation (in log) of absolute values o errors on ellipticity parameters versus the SNR

5 Conclusion

We introduced the SPRITE algorithm which may be seen as an improvement over the super-resolution algorithm used in the software PSFEXTRACTOR. The latter method was implemented iteratively and augmented with a sparse penalty. This

way we achieved significantly better accuracy in terms of ellipticity error.

References

[1] ESA/SRE "Euclid Definition Study Report",2011

[2] S.C. Park, M.K. Park and M.G. Kang, "Super-Resolution Image Reconstruction: a technical overview", IEEE Signal Processing Magazine, **3**:1053–5888, 2003

[3] M. Elab and Y. hel-Or, "A Fast Super-Resolution Reconstruction Algorithm for pure translational motion and common space invariant blur", IEEE Transaction on Image Processing, **10**:1187–1193, 2001

[4] M. Yamagishi, S. Ono and I. Yamada "Two Variants of Alternating Direction Method of Multipliers without Inner Iterations and Their Application to Image Super-Resolution", IEEE International Conference on Acoustics, Speech and Signal Processing, 2012, Kyoto, Japan

[5] J. Fadili and J.-L. Starck, "Monotone operator splitting for fast sparse solutions of inverse problems", SIAM Journal on Imaging and Sciences, 2009, Submitted

[6] H.H. Bauschke, R.S. Burachik, P.L. Combettes, et al. "Fixed-Point Algorithms for Inverse Problems in Science and Engineering" Springer, 185–212, 2011

[7] J.-L. Starck, M. Elad and D. Donoho, "Redundant multiscale transforms and their application for morphological components analysis", Advances in imaging and Electron Physics, **132**:288–348, 2004

[8] A. Beck and M. Teboulle, "A Fast Iterative Shrinkage-Thresholding Algorithm for Linear Inverse Problems", SIAM Journal on Imaging and Sciences,**2**(1):183–202, 2009

[9] E.J. Candès, M.B. Wakin and S.P. Boyd, "Enhancing Sparsity by Reweighted l_1 Minimization", Journal of Fourier Analysis and Applications,**14**:877–905, 2008

[10] M. Cropper et al., "Defining a weak lensing experiment in space", arXiv:astro-ph.IM/1210.7691, 2013

Total Variation Reconstruction From Quasi-Random Samples

Colas Schretter[1,2], Ignace Loris[3], Ann Dooms[1,2] and Peter Schelkens[1,2]

[1]Vrije Universiteit Brussel, Dept. of Electronics and Informatics, Pleinlaan 2, B-1050 Brussels, Belgium

[2]iMinds, Dept. of Multimedia Technologies, Gaston Crommenlaan 8 box 102, B-9050 Ghent, Belgium

[3]Département de Mathématique, Université Libre de Bruxelles, Brussels, Belgium

***Abstract*— Pseudo-random numbers are often used for generating incoherent uniformly distributed sample distributions. However randomness is a sufficient – not necessary – condition to ensure incoherence. If one wants to reconstruct an image from few samples, choosing a globally optimized set of evenly distributed points could capture the visual content more efficiently. This work compares classical random sampling with a simple construction based on properties of the fractional Golden ratio sequence and the Hilbert space filling curve. Images are then reconstructed using a total variation prior. Results show improvements in terms of peak signal to noise ratio over pseudo-random sampling.**

1 Introduction

In compressed sensing (CS) [1], random linear measurements are used to efficiently recover an unknown but assumed sparse signal. We propose the use of quasi-random samples as a better way of sampling. As an initial illustration of the benefits of non-random sampling methods, we reconstruct the Lena image from quasi-random point samples and compare it with the usual pseudo-random point sample reconstruction. As a regularization of the inversion, we use the total variation (TV) prior, *i.e.*, we impose sparsity of the local image gradient.

2 Quasi-random sampling

The law of large number says that any sequence of randomly distributed numbers converge to a uniform distribution. This assumption is however valid only for very large numbers. When a strictly limited budget of samples should be selected for image reconstruction, it might be advantageous to rely on deterministic constructions ensuring more local uniformity in the point distribution, in contrast to clusters of points generated by true (or computer-generated pseudo-) randomness [2].

Elements of the fractional golden ratio sequence $G_s(i)$ with given seed constant $s \in [0,1)$ are given by the fractional part of the sum between s and an integer multiple of the golden ratio:

$$G_s(i) = \{s + i \cdot \phi\}, \forall\, i \geq 1\ , \quad \text{with} \quad \phi = \frac{1+\sqrt{5}}{2} \qquad (1)$$

where $\{t\}$ is the fractional part of the real number t. Note that the conjugate golden section $\tau = \frac{1}{\phi} = \phi - 1$ can be subtitued to ϕ since only fractional parts are retained. Hence, computing $G_s(i+1)$ given $G_s(i)$ only requires an addition and a bit mask.

Such sequences are especially interesting for distributing coordinates on the 1D unit range. A key corollary of the strong irrationality of ϕ (and τ) is that the fractional parts of integer multiples will not align on any regular grid. In previous work [3], the coordinates of this one-dimensional sequence are mapped to higher dimensions, using the invertible Hilbert space filling curve mapping $H : \mathbb{R} \to \mathbb{R}^2$.

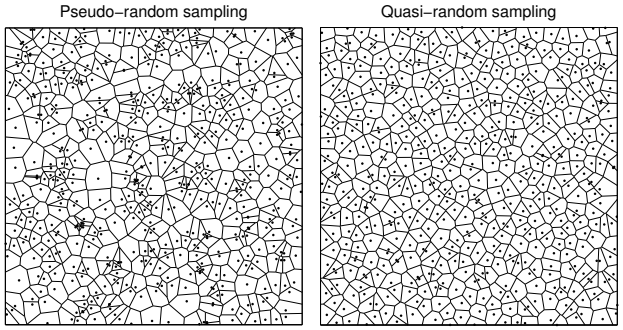

Figure 1: Voronoi diagram for the first 512 points of a pseudo-random and quasi-random sequence. Spurious clusters emerges in the random distribution, while the quasi-random distribution ensures a more evenly uniform coverage.

The Hilbert space filling curve $H\,(t) = (x, y)\,,\ t \in [0, 1)$ defines a nested recursive grid and a locally-preserving traversal order of grid elements. Plugin-in $G_s(\cdot)$ in $H(\cdot)$, we obtain the following sets of N uniformly distributed points:

$$H(G_s(i)) = H(\{s + i \cdot \tau\}),\ i \in [1 \ldots N]. \qquad (2)$$

Figure 1 shows the Voronoi tessellation for two sets of $N = 512$ point samples, using pseudo-random and the proposed quasi-random strategy with $s = 0$. Variability of Voronoi cell's areas indicates increases in local discrepancy.

3 Total variation reconstruction

Let S be the sampling operator $S : \mathbb{R}^{N_1 \times N_2} \to \mathbb{R}^N : u \to y = Su$ that corresponds to choosing a certain N pixels from the $N_1 \times N_2$ image u. Based on samples y of some given image ($y = Su^{\text{orig}}$), the reconstructed image u^{rec} minimizes the total variation and coincides with u^{orig} in the chosen samples, *i.e.*:

$$u^{\text{rec}} = \underset{Su=y}{\arg\min}\ \mathrm{TV}(u). \qquad (3)$$

Here the total variation $\mathrm{TV}(u)$ is defined in terms of the local differences (image gradients) $D(u) \equiv (D_{\text{h}}u, D_{\text{v}}u) \in \mathbb{R}^2$:

$$\begin{cases} D_{\text{h}}u_{i,j} = u_{i+1,j} - u_{i,j} & 1 \leq i \leq N_1 - 1, \\ D_{\text{v}}u_{i,j} = u_{i,j+1} - u_{i,j} & 1 \leq j \leq N_2 - 1, \end{cases} \qquad (4)$$

with $D_{\text{h}}u_{N_1,j} = D_{\text{v}}u_{i,N_2} = 0$. Thus, the total variation $\mathrm{TV}(u)$ in expression (3) is given explicitly by

$$\mathrm{TV}(u) = \sum_{i=1}^{N_1} \sum_{j=1}^{N_2} \sqrt{(D_{\text{h}}u_{i,j})^2 + (D_{\text{v}}u_{i,j})^2}. \qquad (5)$$

I. Loris is a research associate of the Fonds de la Recherche Scientifique-FNRS (Belgium). This work is sponsored by the iMind ICON project 3DUS.

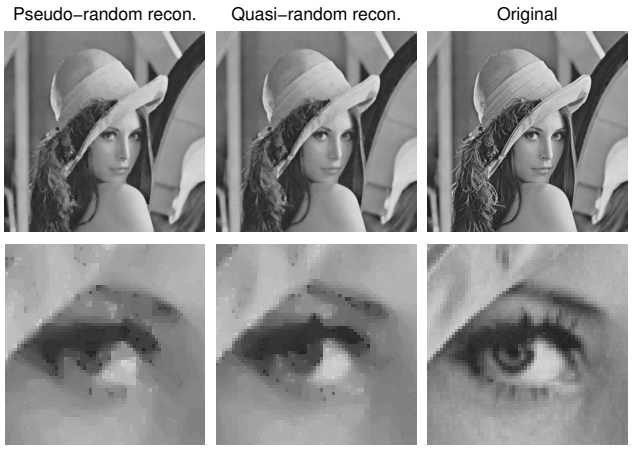

Figure 2: Original Lena image and reconstructions from 50.000 samples using pseudo-random and the proposed quasi-random algorithm. Visual quality and signal to noise ratio are improved when using the quasi-random method.

Total variation was introduced in imaging in 1992[4]. This penalty promotes sparse gradients in the image. In other words, it promotes a piecewise constant reconstruction. In order to compute the solution to the convex minimization problem (3), we use the following iterative scheme:

$$\begin{cases} \bar{u}^{n+1} &= u^n - \tau_1 S^T\left(2v^n - v^{n-1}\right) - \tau_1 D^T w^n \\ w^{n+1} &= P_\lambda\left(w^n + \frac{\tau_2}{\tau_1} D\bar{u}^{n+1}\right) \\ u^{n+1} &= u^n - \tau_1 S^T\left(2v^n - v^{n-1}\right) - \tau_1 D^T w^{n+1} \\ v^{n+1} &= v^n + Su^{n+1} - y \end{cases} \tag{6}$$

where the horizontal (h) and vertical (v) projection P_λ is

$$P_\lambda(w_{\rm h}, w_{\rm v}) = \begin{cases} \lambda\dfrac{(w_{\rm h}, w_{\rm v})}{\sqrt{w_{\rm h}^2 + w_{\rm v}^2}} & \text{if } \sqrt{w_{\rm h}^2 + w_{\rm v}^2} \geq \lambda \\ (w_{\rm h}, w_{\rm v}) & \text{otherwise} \end{cases} \tag{7}$$

and it is understood that it is applied in each pixel separately.

The algorithm (6) converges for positive step sizes $\tau_1 < \|S\|_2^{-2}$ and $\tau_2 < \|D\|_2^{-2}$. In this case, the spectral norm of S is $\|S\|_2 = 1$ and the spectral norm of D is $\|D\|_2 = 2\sqrt{2}$. The convergence speed of this algorithm depends on the positive parameter λ, but its choice is not critical. The algorithm depend on three auxiliary variables ($\bar{u}, v$ and w). All variables are initialized ($n = 0$) by zero. A proof of convergence can be found in [6]. Other methods can be found *e.g.* in [5].

4 Results

In order to compare the two sampling methods, we constructed a numerical experiment in which an image is sampled at a limited number of pixels. The whole image is then reconstructed using the TV penalty.

Figure 2, first row, shows the results of an initial comparison of the proposed quasi-random sampling method and standard pseudo-random sampling. Both reconstructions are based on 50.000 distinct point samples of the Lena image. In both cases we used the TV prior to regularize the inversion from 50.000 samples to a dense image of 512×512 pixels. Figure 2, second row, represents a close-up of the the first row images. More visual detail is discernable in the reconstruction from quasi-random samples, than from pseudo-random samples.

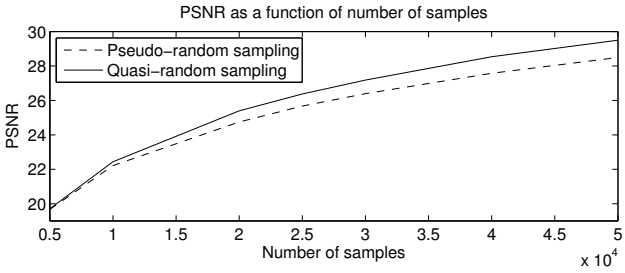

Figure 3: Peak signal to noise ratio (PSNR) is monotonously increasing with the number of images samples available for image reconstruction. The quasi-random sampling strategy yields better recovery of the original image content.

The algorithm was run with up to 1000 iterations. The final peak signal to noise ratio (PSNR) of the reconstruction from quasi-random samples is 29.5dB whereas the final PSNR for the reconstruction from pseudo-random samples is 28.5dB. Furthermore, seven experiments were performed with $5000, 10.000, 20.000, 25.000, 30.000, 40.000$ and 50.000 samples. Figure 3 plots the PSNR for the seven pairs of reconstructions of the Lena image. In both case the PSNR increases monotonically with the number of samples taken, but the quasi-random samples give better reconstruction results in all cases.

5 Conclusion

The law of large numbers says that randomly distributed points are expected to tend towards a uniform distribution. Experiments demonstrate that if the number of point samples is limited, spurious random clustering impairs sensing efficiency since some regions of the sampling domain are more sparsely populated. Fortunately, deterministic quasi-random methods exist for providing evenly distributed but still incoherent points.

Our first results on a natural image were encouraging and future work will quantify the potential of the method on a practical medical imaging setup. Another possibility consists of sampling the wavelet or Fourier domains in quasi-random points, as opposed to the usual pseudo-randomly placed samples, as is often done in applications of the compressed sensing framework.

References

[1] S. Foucart and H. Rauhut, "A mathematical introduction to compressive sensing", Springer, 2013.

[2] H. Niederreiter, "Random number generation and quasi-Monte Carlo methods", CBMS-NSF Regional Conference Series in Applied Mathematics 63, SIAM, Philadelphia, 1992.

[3] C. Schretter and H. Niederreiter, "A direct inversion method for non-uniform quasi-random point sequences", Monte Carlo Methods and Applications, 19(1):1–9, 2013.

[4] L. I. Rudin, S. Osher and E. Fatemi, "Nonlinear total variation based noise removal algorithms", Physica D: Nonlinear Phenomena 60(1–4):259–268, 1992.

[5] A. Chambolle and T. Pock, "A first-order primal-dual algorithm for convex problems with applications to imaging", J. Math. Imaging Vis. 40:120–145, 2011.

[6] I. Loris and C. Verhoeven, "Iterative algorithms for total variation-like reconstructions in seismic tomography", International Journal on Geomathematics, 3(2):179–208, 2012.

Model Selection with Piecewise Regular Gauges

Samuel Vaiter[1], Mohammad Golbabaee[1], Jalal Fadili[2] and Gabriel Peyré[1].
[1]CNRS, CEREMADE, Paris-Dauphine. [2]GREYC, ENSICAEN.

***Abstract*— Regularization plays a pivotal role when facing the challenge of solving ill-posed inverse problems, where the number of observations is smaller than the ambient dimension of the object to be estimated. In such settings, the general approach is to solve a regularized optimization problem, which combines a data fidelity term and some regularization penalty that promotes the low-dimensional/simple structure underlying the vector to recover. This work provides a general framework to capture this low-dimensional structure through what we coin piecewise regular gauges which are convex, non-negative, closed, bounded and positively homogeneous functions. This class of regularizers encompasses many popular examples such as the ℓ^1 norm, $\ell^1-\ell^2$ norm (group sparsity), nuclear norm, as well as several others including the ℓ^∞ norm or the total variation semi-norm. Our main result presents a unified sharp analysis of exact and robust recovery of the low-dimensional subspace model associated to the object to recover from partial measurements.**[1]

1 Regularization by Gauges

We consider the forward model $y = \Phi x_0 + w$, where $y \in \mathbb{R}^Q$ is the vector of observations, $w \in \mathbb{R}^Q$ stands for the noise, and $\Phi \in \mathbb{R}^{Q\times N}$ is a linear operator which maps the N-dimensional signal domain onto the Q-dimensional observation domain. We here consider solutions to the regularized optimization problem

$$x^\star \in \operatorname*{Argmin}_{x\in\mathbb{R}^N} \frac{1}{2}\|y-\Phi x\|^2 + \lambda J(x), \qquad (\mathcal{P}_\lambda(y))$$

or its noiseless counterpart, when $w = 0$,

$$x^\star \in \operatorname*{Argmin}_{x\in\mathbb{R}^N} J(x) \quad \text{subject to} \quad \Phi x = y. \qquad (\mathcal{P}_0(y))$$

Here, J is a convex, non-negative, continuous (hence bounded) and positively homogeneous function. This class of regularizers J include many well-studied ones in the literature. Among them, one can think of the ℓ^1 norm used to enforce sparse solutions [10], the discrete total variation semi-norm [9], the $\ell^1-\ell^2$ norm to induce block/group sparsity [14], the nuclear norm for low-rank matrices [5, 4], or bounded polyhedral gauges [13].

Assuming that the gauges enjoy some piecewise regularity properties, our goal in this paper is to provide a unified analysis of exact and robust recovery guarantees of the subspace model underlying the object x_0 by solving $(\mathcal{P}_\lambda(y))$ from the partial measurements y. As a by-product, this entails a control on $\|x^\star - x_0\|$, where $x^\star$ is the unique minimizer of $(\mathcal{P}_\lambda(y))$.

2 Piecewise Regular Gauges

Associated to any bounded gauge, we define the model subspace at a point x as follows.

[1]This work is a shortened version of [11]

Definition 1 (Model Subspace)**.** *For any vector $x \in \mathbb{R}^N$, let $\bar{S}_x = \operatorname{aff}\partial J(x)$ be the affine hull of the subdifferential at of J at x, and e_x the orthogonal projection of 0 onto $\bar{S}_x$. We denote*

$$S_x = \bar{S}_x - e_x = \operatorname{span}(\partial J(x) - e_x) \quad \textit{and} \quad T_x = S_x^\perp.$$

T_x is coined the model subspace *of x associated to J.*

When J is Gâteaux-differentiable at x, i.e. $\partial J(x) = \{\nabla J(x)\}$, $e_x = \nabla J(x)$ and $T_x = \mathbb{R}^N$. On the contrary, when J is non-smooth at x, the dimension of T_x is strictly less that N, and the regularizing gauge J essentially promotes elements living on this model subspace.

Definition 2 (Antipromoting Gauge)**.** *Let $x \in \mathbb{R}^N \setminus \{0\}$ and $f_x \in \operatorname{ri}\partial J(x)$, where* ri *is the relative interior. The* antipromoting gauge *associated to f_x is the gauge $J^\circ_{f_x}$ whose sublevel set at 1 is $\partial J(x) - f_x$.*

Toward the goal of studying robust recovery by solving $(\mathcal{P}_\lambda(y))$, we introduce a subclass of bounded gauges J for which the mappings $x \mapsto e_x$, $x \mapsto \mathrm{P}_{S_x} f_x$ and $x \mapsto J^\circ_{f_x}$ exhibit local regularity.

Definition 3 (Piecewise Regular Gauge)**.** *A bounded gauge J is said to be* piecewise regular *(PRG for short) at $x \in \mathbb{R}^N$, with respect to the mapping $x \mapsto f_x \in \operatorname{ri}\partial J(x)$ and Γ a coercive gauge bounded on $T_x = T = S^\perp$, if there exist three non-negative reals ν_x, μ_x, τ_x and a real ξ_x such that for every $x' \in T$ with $\Gamma(x-x') \leqslant \nu_x$, one has $T_x = T_{x'}$ (i.e. $x' \in \tilde{T}$) and*

$$\Gamma(e_x - e_{x'}) \leqslant \mu_x \Gamma(x-x'). \qquad (1)$$

$$J^\circ_{f_x}(\mathrm{P}_S(f_x - f_{x'})) \leqslant \tau_x \Gamma(x-x'). \qquad (2)$$

$$\sup_{\substack{u\in S\\ u\neq 0}} \frac{J^\circ_{f_{x'}}(u) - J^\circ_{f_x}(u)}{J^\circ_{f_x}(u)} \leqslant \xi_x \Gamma(x-x'). \qquad (3)$$

Several examples are provided in Table 1. It turns out that the set of piecewise regular gauges is closed by summation and precomposition by a linear operator.

Proposition 1. *Let J and J' two PRG and D a linear operator from $\mathbb{R}^P$ to $\mathbb{R}^N$. Then $J + J'$ and $J \circ D^*$ are PRG.*

3 Model Selection and Robustness

For some model subspace T, the restricted injectivity of Φ on T plays a central role in the sequel. This is achieved by imposing that

$$\operatorname{Ker}(\Phi) \cap T = \{0\}. \qquad (\mathcal{C}_T)$$

Beside condition $(\mathcal{C}_T)$ stated above, the following Identifiability Criterion will play a pivotal role.

Definition 4. *For $x \in \mathbb{R}^N$ such that $(\mathcal{C}_T)$ holds, we define the* Identifiability Criterion *at x as*

$$\mathbf{IC}(x) = J^{\circ}_{f_x}(\Phi_S^* \Phi_T^{+,*} e_x - \mathrm{P}_S f_x),$$

where $T = T_x = S_x^{\perp}$.

It turns out that in such a setting, $\mathbf{IC}(x_0) < 1$ is a sufficient condition for identifiability in the noiseless case, without any any other particular assumption on the bounded gauge J, such as piecewise regularity.

Theorem 1. *Let $x_0 \in \mathbb{R}^N$ and $T = T_{x_0}$. We assume that $(\mathcal{C}_T)$ holds and $\mathbf{IC}(x_0) < 1$. Then x_0 is the unique solution of $(\mathcal{P}_0(y))$.*

Our main contribution is the following result.

Theorem 2. *Let $x_0 \in \mathbb{R}^N$ and $T = T_{x_0}$. We suppose that J is a piecewise regular gauge at x_0 with the corresponding parameters $(\Gamma, \nu_{x_0}, \mu_{x_0}, \tau_{x_0}, \xi_{x_0})$. Assume that $(\mathcal{C}_T)$ holds and $\mathbf{IC}(x_0) < 1$. Then there exist positive constants (A_T, B_T) that solely depend on T, and a constant $C(x_0)$ such that if w and λ obey*

$$\frac{A_T}{1 - \mathbf{IC}(x_0)} \|w\| \leqslant \lambda \leqslant \nu_{x_0} \min\left(B_T, C(x_0)\right) \qquad (4)$$

the solution $x^\star$ of $(\mathcal{P}_\lambda(y))$ with the noisy measurements y is unique, and satisfies $T_{x^\star} = T$. Furthermore, one has

$$\|x_0 - x^\star\| = O\Big(\max(\|w\|, \lambda)\Big).$$

This result asserts that exact recovery of T_{x_0} from noisy partial measurements is possible with the proviso that the regularization parameter λ lies in the interval (4). The value λ should be large enough to reject noise, but small enough to recover the entire subspace T_{x_0}. In order for the constraint (4) to be non-empty, the noise-to-signal level $\|w\|/\nu_{x_0}$ should be small enough, i.e.

$$\frac{\|w\|}{\nu_{x_0}} \leqslant \frac{1 - \mathbf{IC}(x_0)}{A_T} \min\left(B_T, C(x_0)\right) \ .$$

The constant $C(x_0)$ involved in this bound depends on x_0 and has the form

$$C(x_0) = \frac{1 - \mathbf{IC}(x_0)}{\xi_{x_0} \nu_{x_0}} H\left(\frac{D_T\, \mu_{x_0} + \tau_{x_0}}{\xi_{x_0}}\right)$$

where $H(\beta) = \frac{\beta + 1/2}{E_T\, \beta} \varphi\left(\frac{2\beta}{(\beta+1)^2}\right)$ and $\varphi(u) = \sqrt{1+u} - 1$. The constants (D_T, E_T) only depend on T. $C(x_0)$ captures the influence of the parameters $(\mu_{x_0}, \tau_{x_0}, \xi_{x_0})$, where the latter reflect the geometry of the regularizing gauge J at x_0. More precisely, the larger $C(x_0)$, the more tolerant the recovery is to noise. Thus favorable regularizers are those where $C(x_0)$ is large.

4 Relation to Previous Works

This result extends the work of [6] for ℓ^1-regularization, [12] for analysis-ℓ^1, [1] for non-overlapping group Lasso, [2] for the trace norm, and [13] for general polyhedral regularizations. In [3], the authors introduced the notion of decomposable norms. In fact, one can show that their regularizers are a very special subclass of ours. Typical examples of (strongly) decomposable norms are the ℓ^1, $\ell^1 - \ell^2$ and nuclear norms. However, strong decomposability such as ℓ^∞ norm among many others examples, do not fall within their class. The analysis provided in [3] deals only with identifiability in the noiseless case. Their work was extended in [8] when J is the sum of decomposable norms. A notion of decomposability closely related to that of [3], but different, was proposed in [7]. There, the authors study ℓ_2-stability for this class of decomposable norms with a general sufficiently smooth data fidelity. This work however only handles norms, and their stability results require however stronger assumptions than ours.

References

[1] F. Bach. Consistency of the group lasso and multiple kernel learning. *The Journal of Machine Learning Research*, 9:1179–1225, 2008.

[2] F. Bach. Consistency of trace norm minimization. *The Journal of Machine Learning Research*, 9:1019–1048, 2008.

[3] E. Candès and B. Recht. Simple bounds for recovering low-complexity models. *Mathematical Programming*, pages 1–13, 2012.

[4] E.J. Candès and B. Recht. Exact matrix completion via convex optimization. *Foundations of Computational mathematics*, 9(6):717–772, 2009.

[5] M. Fazel. *Matrix Rank Minimization with Applications. PhD thesis*. PhD thesis, Stanford University,, 2002.

[6] J.-J. Fuchs. On sparse representations in arbitrary redundant bases. *IEEE Transactions on Information Theory*, 50(6):1341–1344, 2004.

[7] S. Negahban, P. Ravikumar, M.J. Wainwright, and B. Yu. A unified framework for high-dimensional analysis of m-estimators with decomposable regularizers. *arXiv preprint arXiv:1010.2731*, 2010.

[8] S. Oymak, A. Jalali, M. Fazel, Y.C. Eldar, and B. Hassibi. Simultaneously structured models with application to sparse and low-rank matrices. *arXiv preprint arXiv:1212.3753*, 2012.

[9] L.I. Rudin, S. Osher, and E. Fatemi. Nonlinear total variation based noise removal algorithms. *Physica D: Nonlinear Phenomena*, 60(1):259–268, 1992.

[10] R. Tibshirani. Regression shrinkage and selection via the Lasso. *Journal of the Royal Statistical Society. Series B. Methodological*, 58(1):267–288, 1996.

[11] S. Vaiter, M. Golbabaee, J. Fadili, and G. Peyré. Model selection with piecewise regular gauges. *arXiv preprint arXiv:1307.2342*, 2013.

[12] S. Vaiter, G. Peyré, C. Dossal, and M. J. Fadili. Robust sparse analysis regularization. *IEEE Transactions on Information Theory*, 59(4):2001–2016, 2013.

[13] S. Vaiter, G. Peyré, and M. J. Fadili. Robust polyhedral regularization. In *Proc. SampTA*, 2013.

[14] M. Yuan and Y. Lin. Model selection and estimation in regression with grouped variables. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 68(1):49–67, 2005.

J	T_x	e_x	$J^{\circ}_{f_x}$ on S_x	ν_x	μ_x	Comment
$\|\cdot\|_1$	$\{\eta :\ \forall j \notin I,\ \eta_j = 0\}$	$\mathrm{sign}(x)$	$\|\cdot\|_\infty$	$\min_{i \in I} \|x_i\|$	0	$I = \mathrm{supp}(x)$
$\|D^*\cdot\|_1$	$\mathrm{Ker}(D^*_{I^c})$	$\mathrm{P}_{\mathrm{Ker}(D^*_{I^c})} \mathrm{sign}(D^* x)$	$\inf_{z \in \mathrm{Ker}(D_{I^c})} \|D^+_{I^c}\eta + z\|_\infty$	$\min_{i \in I} \|\langle d_i,\ x_i\rangle\|$	0	$I = \mathrm{supp}(D^* x)$
$\|\cdot\|_{1,2}$	$\{\eta :\ \forall j \notin I,\ \eta_j = 0\}$	$(\mathcal{N}(x_b))_{b \in \mathcal{B}}$	$\|\cdot\|_{\infty,2}$	$\min_{b \in I} \|x_b\|$	$\frac{\sqrt{2}}{\nu_x}$	$I = \{g \in \mathcal{B} :\ x_g \neq 0\}$
$\|\cdot\|_*$	$\{Z :\ U_\perp^* Z V_\perp = 0\}$	UV^*	$\|\cdot\|_{sp}$	$\frac{1}{4}\min_{i \in I} \|\sigma_i\|$	$\frac{1}{\nu_x}$	$x = U\Lambda V^*$
$\|\cdot\|_\infty$	$\{\alpha :\ \alpha_I = \rho s_I \text{ for } \rho \in \mathbb{R}\}$	$\mathrm{sign}(x)/\|I\|$	$\max_{i \in I} (-\|I\| s_i \eta_i)_+$	$(\|x\|_\infty - \max_{j \notin I} \|x_j\|)$	0	$I = \{i :\ \|x_i\| = \|x\|_\infty\}$

Table 1: Examples of Piecewise Regular Gauges. For all these regularizations, $\tau_x = \xi_x = 0$.

Adaptive Compressed Sensing Using Sparse Measurement Matrices

Dejan Vukobratović[1] and Aleksandra Pižurica[2]

[1]Department of Power, Electronics and Communications Engineering, University of Novi Sad, Serbia.

[2]Department of Telecommunications and Information Processing, Ghent University, Belgium.

***Abstract*— Compressed sensing methods using sparse measurement matrices and iterative message-passing recovery procedures are recently investigated due to their low computational complexity and excellent performance. The design and analysis of this class of methods is inspired by a large volume of work on sparse-graph codes such as Low-Density Parity-Check (LDPC) codes and the iterative Belief-Propagation (BP) decoding algorithms. In particular, we focus on a class of compressed sensing methods emerging from the Sudocodes scheme that follow similar ideas used in a class of sparse-graph codes called rateless codes. We are interested in the design and analysis of adaptive Sudocodes methods and this paper provides initial steps in this direction.**

1 Introduction

Let $\mathbf{x} = (x_1, x_2, \ldots, x_N) \in \mathbb{F}_q^N$ be a vector of input symbols (message) from q-ary alphabet. Let us transform this message into arbitrarily long sequence of output symbols $\mathbf{y} = (y_1, y_2, \ldots) \in \mathbb{F}_q$, where each output symbol $y_i \in \mathbb{F}_q$ is a "projection" of message $\mathbf{x}$ onto a sparse binary vector $\mathbf{g}_i$ of length N, i.e., $y_i = \mathbf{g}_i \cdot \mathbf{x}^T$. Each vector $\mathbf{g}_i$ is obtained simply by distributing small number d of ones uniformly at random over N possible positions. We assume the number of ones in $\mathbf{g}_i$, called the degree of $\mathbf{g}_i$, is drawn from a degree distribution Ω over integers $d \in [1, 2, \ldots, N]$. LT codes are class of rateless codes providing a recipe for capacity-approaching degree distributions Ω that, combined with a simple iterative recovery procedure, recover $\mathbf{x}$ from any $N + O(\sqrt{N}\ln^2(N/\delta))$ output symbols with probability $1 - \delta$ [1]. The average degree of Ω scales as $O(\ln(N/\delta))$ and thus results in transformation/recovery complexity of $O(N\ln(N/\delta))$.

If $\mathbf{x} \in \mathbb{R}^N$ and is K-sparse (meaning that only $K << N$ input symbols are non-zero) the above rateless methodology for message recovery is still applicable. Sudocodes scheme [2], that independently emerged in compressed sensing community, exploited similar ideas and proposed the iterative recovery procedure which is later identified to be an instance of the verification-decoding for sparse-graph codes [3]. The Sudocodes scheme is extended into a more general framework called Compressed Sensing via Belief Propagation (CSBP) inspired by the iterative Belief-Propagation (BP) decoding of LDPC codes [4]. Sudocodes sparked significant interest of coding community for Compressed Sensing (CS) methods using sparse measurement matrices and iterative message-passing recovery algorithms. Using well-developed tools from coding theory, CS methods combined with verification-based recovery have been analyzed in [5] and [6].

This paper explores Sudocodes scheme using design and analysis tools used in rateless coding. These tools are applied for the design of *adaptive* Sudocodes scheme, where the sequence of output symbols are not mutually independent but designed based on the outcomes of previous output symbols.

2 The Sudocodes: Design and Analysis

The Sudocodes scheme projects the K-sparse message $\mathbf{x} \in \mathbb{R}^N$ into a sequence of output symbols $y_i \in \mathbb{R}$ using random sparse binary projection vectors $\mathbf{g}_i$ of constant degree $d = L$. In other words, the Sudocodes scheme is a rateless code applying the degree distribution[1] $\Omega(x) = x^L$.

The Sudocode scheme recovers $\mathbf{x}$ from a sequence of M output symbols $\mathbf{y} = (y_1, y_2, \ldots, y_M)$ by applying verification-based recovery across a measurement graph [3]. A measurement graph consists of N input symbol nodes corresponding to $\mathbf{x}$ and M output symbol nodes corresponding to $\mathbf{y}$. Edges of the graph connect each output symbol node y_i to its neighbor set $N(y_i)$ of input symbol nodes determined by non-zero positions in $\mathbf{g}_i$. The measurement graph is usually defined by (edge-oriented) degree distributions for input and output symbol nodes: $\lambda(x) = \sum_i \lambda_i \cdot x^{i-1}$ and $\omega(x) = \sum_i \omega_i \cdot x^{i-1} = \Omega'(x)/\Omega'(1)$ [7]. Two verification-based recovery algorithms called LM1 and LM2 algorithm are proposed in [3]:

LM1: The LM1 operates iteratively over the measurement graph by applying following rules: 1) If $y_i = 0$ then $\forall x_j \in N(y_i) : x_j = 0$; Verify all $x_j : x_j \in N(y_i)$. 2) If $(y_i \neq 0) \wedge (|N(y_i)| = 1)$ then $x_j = y_i$ for the node $x_j \in N(y_i)$; Verify x_j. 3) Remove verified coefficient nodes and their incident edges from the graph; Subtract out verified values from remaining measurements. 4) Repeat until successful signal recovery or makes no progress in two consecutive iterations.

LM2: Besides the above LM1 rules, LM2 adds the additional one: If $(N(y_i) \cap N(y_j) = \{x_k\}) \wedge (y_i = y_j)$ then $x_k = y_i = y_j$ and $\forall x_l \in \{N(y_i) \cup N(y_i) \setminus x_k\} : x_l = 0$; Verify all $x_l \in \{N(y_i) \cup N(y_i) \setminus x_k\}$.

Reconstruction of $\mathbf{x}$ using LM1 algorithm is equivalent to LT decoding [1], thus the AND-OR-tree analysis [7] used to asymptotically evaluate the performance of LT codes can be easily reshaped for LM1 recovery analysis.

Lemma 2.1. *Let $p_l^{(z)}$ and $p_l^{(\bar{z})}$ denote the probabilities that a zero and non zero input symbol, respectively, is not recovered after l iterations of LM1 recovery. Then*

$$p_l^{(\bar{z})} = \lambda\Big(1 - \sum_{i=1}^{d_{max}^{(m)}} \omega_i \cdot \sum_{j=0}^{i-1} \binom{i-1}{j} (\alpha \bar{p}_{l-1}^{(\bar{z})})^j \cdot (\bar{\alpha}\bar{p}_{l-1}^{(z)})^{i-1-j}\Big)$$

$$p_l^{(z)} = \lambda\Big(1 - \sum_{i=1}^{d_{max}^{(m)}} \omega_i \cdot \sum_{j=0}^{i-1} \binom{i-1}{j} (\alpha \bar{p}_{l-1}^{(\bar{z})})^j \cdot \bar{\alpha}^{i-1-j}\Big),$$

where $\alpha = K/N$, we use compact notation $\bar{x} = 1 - x$, and the recursion is initialized at $p_0^{(z)} = p_0^{(\bar{z})} = 1$. Finally, $p_l = \alpha p_l^{(\bar{z})} + \bar{\alpha} p_l^{(z)}$ is the average probability that a signal coefficient is not recovered after l iterations of LM1 recovery.

[1]In rateless coding, the degree distributions are represented as polynomials $\Omega(x) = \sum_i \Omega_i \cdot x^i$, where $\Omega_i = \mathbb{P}(d = i)$.

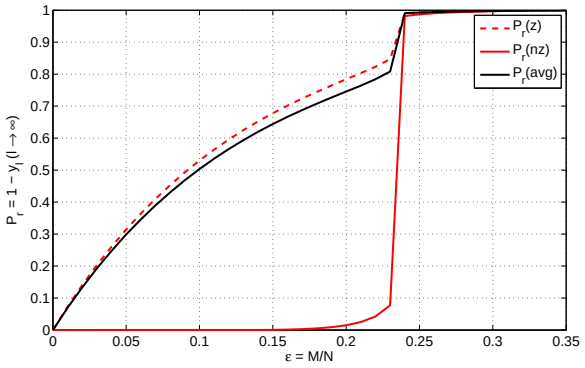

Figure 1: Asymptotic performance of Sudocodes scheme.

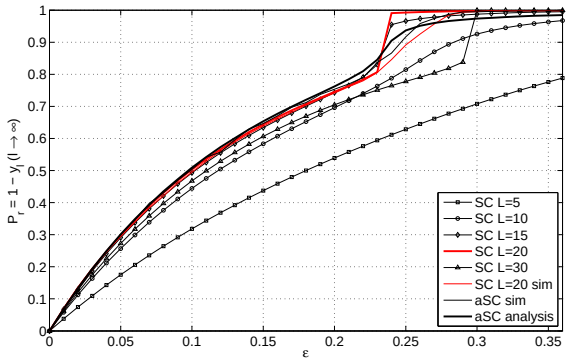

Figure 3: Sudocodes (SC) and adaptive-SC (aSC) performance.

In terms of zero input symbol recovery probability $\mathbb{P}_r^{(z)} = 1 - p^{(z)}$, it is easy to show that the optimal degree distribution equals $\Omega(x) = x^{d^*}$, where $d^* \approx (\log \frac{1}{1-\alpha})^{-1}$. Optimal degree distributions for average input symbol recovery probability $\mathbb{P}_r = 1 - p$ can be obtained as described in [1] and [8].

Example 2.1. *Fig. 1 shows asymptotic recovery probabilities (as $N \to \infty$) obtained from* Lemma 1 *for Sudocodes scheme that applies $\Omega(x) = x^{20}$ after LM1 recovery of input message of sparsity-factor $\alpha = 0.05$ ($d^*(\alpha = 0.05) = 20$) for zero and non-zero input symbols and average value. Fig. 3 shows asymptotic recovery probability curves for $\Omega(x) = x^d$ with increasing $d = \{5, 10, 15, 20, 25, 30\}$.*

3 Adaptive Sudocodes

Consider the following modification to the Sudocodes scheme:

Modification 1: After generating output symbol $y_i = 0$, remove input symbols $x_j \in N(y_i)$ from consideration while generating following output symbols.

Modification 2: Due to the first modification, the sparsity-factor α decreases as the number of generated output symbols increase. Thus we continuously update the optimal degree $d^* = d^*(\alpha)$ before each output symbol is generated.

We analyze the Adaptive Sudocodes by following the evolution of measurement matrix with the process of generating output symbols (see Fig. 2). Starting from the initial sparsity-factor $\alpha_0 = K/N$, the process runs through a sequence of increasing α-values $\{\alpha_1, \alpha_2, \ldots\}$ at which the optimal d^*-values decrement. The total of m_i measurements are generated using degree d_i^* during which the sparsity-factor increases from α_i to α_{i+1}. By tracking the parameters (dimensions and degree distributions) of the sequence of measurement matrices $\mathbf{\Phi}_i, i \geq 0$, and using *Lemma 1*, we are able to track the evolution of asymptotic recovery probabilities of LM1 reconstruction. The complete analysis is omitted due to lack of space, however, we note that due to approximations of degree distributions, the resulting analysis is not exact (see example below).

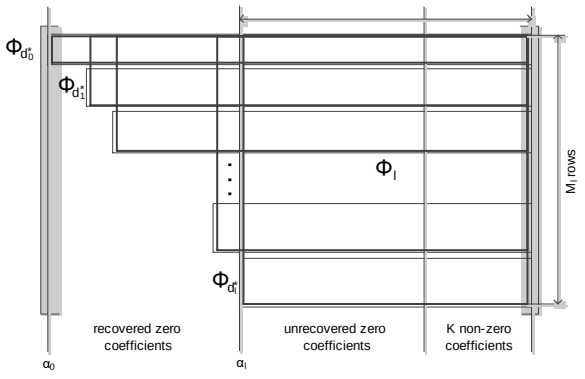

Figure 2: Evolution of measurement matrix in Adaptive Sudocodes.

Example 3.1. *Fig. 3 shows asymptotic recovery probability for Adaptive Sudocodes for initial sparsity-factor $\alpha_0 = 0.05$. Simulated results for both Sudocodes and Adaptive Sudocodes scheme are provided for $K = 50$ and $N = 1000$ demonstrating that Adaptive Sudocodes performs better with lower complexity since the optimal degree d^* decreases from $d = 20$ down to $d = 2$ at $M/N = 0.35$. We note that our approximate analysis matches well the simulated performance except that it becomes conservative in the region of large M/N-values.*

4 Comments and Future Work

The rateless codes that apply $\Omega(x)$ with constant average degree $\mu = \Omega'(1)$ are affected by upper bound on recovery probability that scales as $e^{-\mu \frac{M}{N}}$. In LT codes, this is solved by using degree distributions whose average degree scales as $O(\log N)$ [1], for the price of increased encoding/decoding complexity of $O(N \log N)$. More advanced solution called Raptor codes preserves linear encoding/decoding complexity by using high-rate precoding combined with the constant average degree $\Omega(x)$ [8].

The Sudocodes scheme suffers from the same error-floor problem due to a constant-degree $\Omega(x)$ employed. This is solved by employing second phase where non-sparse measurements are combined with matrix inversion to recover small fraction of remaining input symbols [2]. However, instead of "postcoding", it is more instructive to use the Raptor-idea of "precoding" the message by adding small number of additional precoded input symbols followed by Sudocodes scheme with constant-degree $\Omega(x)$. This design is part of our ongoing work.

Finally, extending Adaptive Sudocodes to LM2 recovery has strong potential for further performance improvements and knowledge extraction from previous measurements. Note that in LM1 recovery, zero input symbols are learned only from zero-valued output symbols, while in LM2 recovery, the additional recovery rule provides new possibilities for zero input symbols identification. However, asymptotic analysis of LM2 algorithm seems to be considerably more involved [6]. The design and analysis of Adaptive Sudocodes for LM2 recovery is part of our ongoing work.

References

[1] M. Luby: "LT Codes," Proc. IEEE Symp. Found. of Comp. Science (FOCS), Vancouver, Canada, Nov. 2002.

[2] S. Sarvotham, D. Baron and R. G. Baraniuk: "Sudocodes: Fast Measurement and Reconstruction of Sparse Signals," *IEEE Int'l Symp. Information Theory - ISIT 2006*, pp. 2804–2808, Seattle, USA, July 2006.

[3] M. Luby and M. Mitzenmacher:"Verification-based decoding for packet-based low-density parity-check codes," *IEEE Trans. Information Theory*, Vol 51, No 1, pp. 120–127, January 2005.

[4] D. Baron, S., Sarvotham, R. G. Baraniuk: "Bayesian compressive sensing via belief propagation," *IEEE Trans. Signal Processing,* Vol. 58, No. 1, pp. 269–280, January 2010.

[5] F. Zhang and H. D. Pfister: "Verification decoding of high-rate LDPC codes with applications in compressed sensing," *IEEE Trans. Information Theory,* Vol. 58, No. 8, pp. 5042–5058, August 2012.

[6] Y. Eftekhari, A. Heidarzadeh, A. H. Banihashemi and I. Lambadaris: "Density Evolution Analysis of Node-Based Verification-Based Algorithms in Compressed Sensing," *IEEE Trans. Information Theory,* Vol. 58, No. 10, pp. 6616–6645, October 2012.

[7] M. Luby, M. Mitzenmacher and A. Shokrollahi: "Analysis of random processes via and-or tree evaluation," *ACM Symp. on Discrete algorithms - SODA 1998,* pp. 364–373, San Francisco, USA, Jan. 1998.

[8] A. Shokrollahi: "Raptor codes," *IEEE Trans. Info. Theory,* Vol. 52, No. 6, pp. 2551–2567, June 2006.

Finished on October 10, 2014, Louvain-la-Neuve, Belgium.